The current issue and full text archive of this journal is available at
www.emeraldinsight.com/1756-378X.htm

Foundations of swarm robotic chemical plume tracing from a fluid dynamics perspective

Diana F. Spears
Swarmotics, LLC, Laramie, Wyoming, USA
David R. Thayer
Physics and Astronomy Department, University of Wyoming, Laramie, Wyoming, USA, and
Dimitri V. Zarzhitsky
Computer Science Department, University of Wyoming, Laramie, Wyoming, USA

Received 6 February 2009
Revised 22 July 2009
Accepted 15 August 2009

Abstract

Purpose – In light of the current international concerns with security and terrorism, interest is increasing on the topic of using robot swarms to locate the source of chemical hazards. The purpose of this paper is to place this task, called chemical plume tracing (CPT), in the context of fluid dynamics.

Design/methodology/approach – This paper provides a foundation for CPT based on the physics of fluid dynamics. The theoretical approach is founded upon source localization using the divergence theorem of vector calculus, and the fundamental underlying notion of the divergence of the chemical mass flux. A CPT algorithm called fluxotaxis is presented that follows the gradient of this mass flux to locate a chemical source emitter.

Findings – Theoretical results are presented confirming that fluxotaxis will guide a robot swarm toward chemical sources, and away from misleading chemical sinks. Complementary empirical results demonstrate that in simulation, a swarm of fluxotaxis-guided mobile robots rapidly converges on a source emitter despite obstacles, realistic vehicle constraints, and flow regimes ranging from laminar to turbulent. Fluxotaxis outperforms the two leading competitors, and the theoretical results are confirmed experimentally. Furthermore, initial experiments on real robots show promise for CPT in relatively uncontrolled indoor environments.

Practical implications – A physics-based approach is shown to be a viable alternative to existing mainly biomimetic approaches to CPT. It has the advantage of being analyzable using standard physics analysis methods.

Originality/value – The fluxotaxis algorithm for CPT is shown to be "correct" in the sense that it is guaranteed to point toward a true source emitter and not be fooled by fluid sinks. It is experimentally (in simulation), and in one case also theoretically, shown to be superior to its leading competitors at finding a source emitter in a wide variety of challenging realistic environments.

Keywords Robotics, Chemical technology, Fluid dynamics

Paper type Research paper

International Journal of Intelligent Computing and Cybernetics
Vol. 2 No. 4, 2009
pp. 745-785
© Emerald Group Publishing Limited
1756-378X
DOI 10.1108/17563780911005863

This paper was partially supported by the National Science Foundation, Grant No. NSF44288.

1. Introduction

The chemical plume tracing (CPT) task is a search and localization problem, in which one must find locations where potentially harmful substances are being released into the environment. The ability to quickly locate chemical sources is an integral component of numerous manufacturing and military activities. In light of the current international concerns with security and the possibility of a chemical terrorist attack, many private and government agencies, including the Japan Defense Agency, the US Department of Homeland Security, the US Defense Advanced Research Projects Agency (DARPA), the Danish Environmental Protection Agency, the UK Health Protection Agency (HPA), and local emergency response agencies throughout the world have expressed interest in updating techniques used to track hazardous plumes, and improving the search strategies used to locate an active chemical toxin emitter Board on Atmospheric Sciences and Climate (2003), Shinoda (2003), and Chemical Hazards and Poisons Division, HPA (2006). As a result, there is a growing interest among both researchers and engineers in designing robot-based approaches to solve this problem.

There are two main motivations for this paper. The first is to cast CPT as a fluid dynamics problem. The second is to clarify the advantages of using a fluid-based CPT approach for solving a fluid-based problem. When the *Environmental Fluid Mechanics* journal published a special issue on the topic of CPT in Cowen (2002), the papers in that issue were groundbreaking in the sense that they laid an initial foundation for the field. Like the majority of prior and current CPT strategies, the strategies presented in the special issue were based on observations of living organisms, e.g. insects, land animals, aerial, and aquatic creatures. In other words, the CPT strategies outlined in that special issue were biomimetic, i.e. they were designed to mimic biological systems. The objective of this paper is to establish a theoretical foundation for an alternative type of strategy – one that is physicomimetic, or physics-based, i.e. it is designed to mimic the way in which physical particles such as molecules behave. Swarm robotic systems based on physicomimetics, also called artificial physics, use virtual forces to maintain a desired inter-robot distance (Spears and Gordon, 1999). Using such virtual forces, the robots can achieve lattice formations, which are ideal for computing derivatives. Virtual forces then attract the lattice toward desirable regions in the environment, such as chemical density maxima. Our new physicomimetic approach to multi-robot CPT exploits fundamental principles of fluid dynamics to drive a swarm lattice to a chemical source emitter.

What motivated us to explore a physics-based alternative? The chief difficulty with biomimetic algorithms is that in order to mimic natural biological strategies with robots, sophisticated sensors are needed. However, despite recent advances in manufacturing techniques, state-of-the-art chemical sensors remain crude and inefficient compared to their biological counterparts. In short, it is difficult to achieve optimal or near-optimal performance in artificial CPT systems by relying on mimicry of biological systems alone. For example, Crimaldi *et al.* (2002) analyzed crustaceans' "antennae" sensor arrays, also called "olfactory appendages," and found that an antennae structure with multiplicity and mobility improves an organism's CPT performance. But we can surpass this with cooperating mobile robots. In particular, the inherent flexibility afforded by the multiplicity and mobility of a swarm (large group) of robots far exceeds that of a single organism's antennae.

Throughout this paper, we will emphasize the advantage of this swarm-based physicomimetic CPT approach, in which autonomous robots employ strategies derived from many years of advanced human study on the topic of fluid dynamics. We are fully aware of the limitations on the ability to predict and model advanced turbulent flow dynamics and chemical transport due to the non-linear and stochastic nature of these processes. However, as we demonstrate in this paper, we already have enough understanding and mathematical expertise at our disposal to design a complete robotic system for locating chemical sources based on the physical cues extracted from local fluid observations. Furthermore, our ability to handle both laminar and turbulent flow regimes (the latter containing filaments or packets of chemical that are not fully diffused) with our CPT approach has been demonstrated experimentally (Zarzhitsky and Spears, 2005; Zarzhitsky *et al.*, 2004a, b, c, 2005; Spears D. *et al.*, 2005) for a variety of different experiments.

If physicomimetic (i.e. physics-based) strategies based on the fundamentals of fluid dynamics can be shown to complement existing search techniques, then autonomous land, marine, and micro-air vehicles can use physicomimetic or hybrid (combined) physicomimetic-biomimetic CPT algorithms to improve their performance. Here, we explore the theoretical underpinnings of plume tracing and source localization strategies that are derived from a fundamental understanding of fluids. Since the emphasis of this paper is on the fluid dynamics foundation of CPT, we examine the most fluid-based approach to CPT, called fluxotaxis, in much greater depth than the other CPT strategies. Theoretical results presented in this paper show that fluxotaxis is competitive with or outperforms the two most popular CPT strategies, under realistic conditions. Furthermore, we have implemented and extensively tested our fluxotaxis approach in simulation and demonstrated its superiority over other CPT algorithms under a wide variety of realistic flow conditions (Zarzhitsky, 2008). Performance metrics included precision and consistency of finding the emitter, and swarm scalability.

Section 2 presents an overview of the CPT problem and its practical importance. The relevance of this work to the design of intelligent robotic systems is examined in Section 3, where we place the fluid analyses in the context of multi-robot CPT; an implementation and experiments with CPT on real hardware robots is discussed. Then, Sections 4 and 5 provide the necessary background for understanding how our CPT approach is a natural application of fluid physics. Section 6 details the specifics of the CPT task and its subtasks, and then Sections 7 and 8 provide an overview of related work on two of these subtasks. In Section 8, we describe the fluxotaxis technique as the first robotic CPT approach based on an application of fluid dynamics in the context of chemical plumes and emitters. Since fluxotaxis is the most fluid-oriented of all the CPT approaches, it is further theoretically analyzed in the remainder of the paper. Section 9 contains a formal analysis of the fluxotaxis approach, and Section 10 compares fluxotaxis with the most widely employed chemotaxis strategy from the perspective of fluid dynamics. Our analytical results provide a quantitative measure of how effective these strategies are in predicting the direction of the source chemical emitter. Section 11 provides experimental evidence in simulation of the practical value of the theory. Finally, Section 12 presents conclusions and future work.

2. CPT: an important and challenging problem

The objective of the CPT task is rapid localization of the emitter that is ejecting the chemical. Plume tracing is different from plume mapping – the objective of the latter is to acquire a map of the plume. In CPT, complete mapping is typically unnecessary – only the source emitter must be found. The underlying motivation for CPT is typically to minimize human exposure to the toxin; therefore, the use of autonomous robots is especially appropriate. It is difficult to predict ahead of time the conditions that the plume-tracing robots will encounter at the site of the contamination – it may be an urban setting after a deliberate terror act, or an industrial facility after a chemical accident (Figure 1). Therefore, it is not surprising that both government and commercial customers have expressed interest in obtaining an effective CPT solution (Caldwell *et al.*, 1997; Cordesman, 2001; Hsu, 2003; Coirier *et al.*, 2005).

Evolution of chemical plumes occurs in three dimensions. However, in most cases, due to practical constraints, CPT activities must be carried out in a narrow, two-dimensional horizontal plane. The CPT problem is even more challenging when the plume-tracing robots occlude each other, or perturb the surrounding flow due to their own movement. Since, in this paper, we analyze the main characteristics of several different plume-tracing algorithms, we ablate local effects of the sensing robots on the flow. It is difficult enough to study CPT without this complication. In fact, this paper omits a discussion of robotic sensing issues, so as to avoid confounding the main presentation of the relevant fluid dynamics. Instead, we focus on the theoretical properties of fluid flow, and what chemical signals and signatures the robots should follow in order to track the plume toward the source, as well as the information necessary to correctly recognize the source emitter. Future articles will address sensing and other hardware issues.

3. Swarm robotics for CPT

The strong emphasis on the physical foundations of fluid flow is due to our interest in a multi-robot approach to the CPT problem. We define a "robotic swarm" as a team of

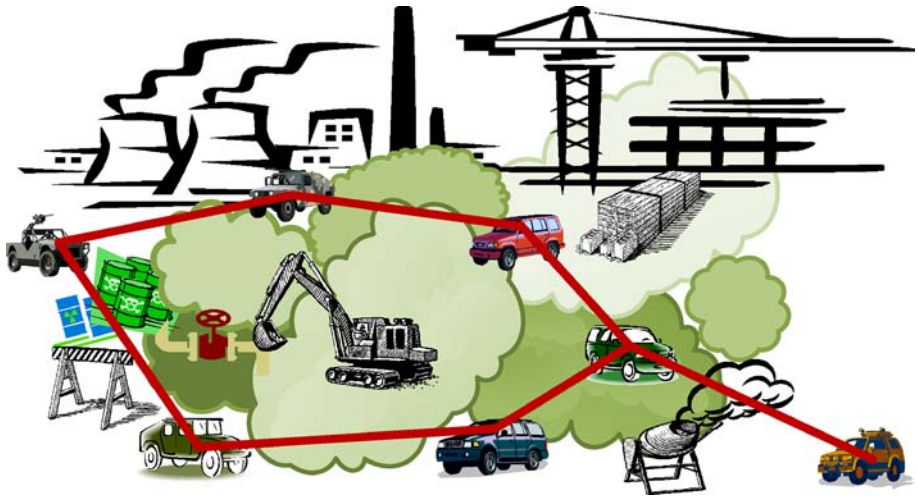

Figure 1.
Envisioned CPT scenario of an intelligent robotic swarm responding to a chemical accident at an industrial facility

Note: The autonomous drones establish an *ad hoc* communication and sensor network (bold lines) to collaboratively search the area for the source of the toxin

mobile autonomous vehicles that cooperatively solve a common problem. The swarm size and type may vary anywhere from a few simple, identical robots to thousands of different hardware platforms, including ground-based, aquatic (autonomous surface/underwater vehicles), or aerial (unmanned air vehicles). In our swarm design, each mobile robot serves as a dynamic node in a distributed sensor and computation network, so that the swarm as a whole functions as an adaptive sensing grid, constantly sharing various fluid measurements between adjacent vehicles. Weissburg *et al.* (2002) provide convincing evidence for the CPT performance advantage of sampling the plume at multiple, spatially separated locations simultaneously. Spears *et al.* (2004) have confirmed this finding in numerous experiments on a group of autonomous CPT robots.

We want to emphasize the importance and value of consistent adherence to the physicomimetic design philosophy for each component of the system. In order to understand the benefits of implementing our physics-based approach, consider the following three functional aspects of a swarm of physicomimetic robots. The most important facet is of course, the robots' cooperative behavior, which we achieve by emulating real-world physics in the on-board control software, using virtual "social" forces for maintaining desired inter-robot distances. In other words, robots that use physicomimetics act as particles, analogous to molecules within a solid, liquid, or gas, so that the whole swarm mimics properties of a real-world physical substance. This enables us to use standard physics analysis techniques to estimate the swarm's performance prior to its deployment (Spears W. *et al.*, 2005). The resulting ability to analytically predict both the short- and long-term behavior of physicomimetic systems offers tremendous advantages for fielding reliable robotic collectives.

At the same time, we can also view the swarm of intelligent robots as a distributed, adaptive computational fluid dynamics grid or mesh. This functionality is the direct consequence of structured swarm formations, because adjacent vehicles can share their local fluid observations, thus boosting the effective resolution of their on-board flow sensors. Our physicomimetics-driven swarm system is a viable and compelling alternative to static *in situ* sensing grids.

Lastly, observe that the on-going exchange of sensor measurements between vehicles helps the robots to process sensor data simultaneously from many different locations. Therefore, we can employ the swarm as a distributed, mobile computer capable of dynamic configuration and run-time load balancing. This unique feature of physicomimetics allows our robotic systems to address the growing interest in distributed computing with a flexible, low maintenance solution, ideally suited for deployment in areas lacking the necessary communication and data-processing infrastructure. Furthermore, the swarm performs fully distributed computation in the sense that each robot in the swarm shares sensor values with its immediate neighbors only, and decides where to move next based only on that strictly local information. There is no explicit collaboration between distant swarm members. Aggregate swarm behavior emerges implicitly, without any global coordination.

Since our present goal is to test the swarm on the plume tracing task, this paper defines some of the fluid dynamics equations and mathematical approximation techniques that intelligent robots can employ to improve their emitter localization performance. We also consider requirements for implementing and analyzing a physics-based swarm-oriented chemical source localization algorithm called

"fluxotaxis," which we derived from theoretical analyses of fluid-driven chemical transport effects. The fluxotaxis-driven swarm moves like a fluid through the plume, and the robots' on-board software, which is based on fluid theory described here, analyzes chemical emissions and the ambient fluid in order to perform effective chemical source localization. In other words, our CPT approach is physicomimetic in two different respects. First, the robots use physicomimetics to stay in fluid-like (bendable) lattice formations as they trace the plume in search of the chemical emitter. Second, when executing the fluxotaxis CPT algorithm, the robots use physicomimetic methods for analyzing the fluid flow and rely on fluid dynamics principles to estimate the location of the chemical emitter. The remainder of this paper will focus on the latter physicomimetic capabilities of the swarm, i.e. we will describe the CPT algorithms used by the robots, instead of the virtual physics that organizes the robots into formations.

We have demonstrated in a robot-faithful simulation, with swarms of hundreds of robots, that our physicomimetic CPT theory is directly applicable to a distributed swarm robotic paradigm. Fluid velocity and chemical concentration are measured independently by each robot and, as we stated above, the collected sensor data are communicated between neighboring vehicles in order to calculate (partial) derivatives. Derivative calculations are performed by the robots using the numerical finite-difference method of second-order accurate central differences (Zarzhitsky, 2008). Note that this short-range sharing of information between robots is sufficient for calculating local navigation gradients, so there is never a need for a "global broadcast" of these values. In other words, implementation of the theory we present here requires no leaders or central controllers or global information, which in turn increases the system's overall robustness, while lowering its cost.

Zarzhitsky *et al.* (2005) show that fluxotaxis is a practical algorithm that can be implemented on a swarm of inexpensive, distributed, autonomous, collaborating robots. They describe how the basic fluxotaxis algorithm is tested, refined and evaluated in comprehensive software simulations, and introduce swarm-specific CPT performance evaluation metrics. The simulation work is of interest because it solves the multi-robot cooperation and coordination problem by using physicomimetic design based on virtual (artificial) social forces (Spears *et al.*, 2004), as described above. This allows the robots to self-organize in an emergent (i.e. without explicit programming) fashion into a dynamic sensor lattice as they navigate toward the chemical emitter. Spears D. *et al.* (2005) show that this swarm-based distributed computer is well-suited for performing fluid-dynamic computations.

As part of our comprehensive research effort, we applied one of the CPT algorithms discussed in this paper (chemotaxis, which follows the chemical density gradient) to our mobile robot prototype, and then ran tests using laboratory-scale ethanol vapor plume configurations. Figure 2, adapted from Spears *et al.* (2006), shows our distributed CPT

Figure 2.
A laboratory-scale CPT scenario, in which miniature robotic vehicles successfully find a source of ethanol vapor (top-right)

implementation, featuring three robots equipped with a miniature metal-oxide volatile organic compound sensor, in addition to an inexpensive 16-bit microcontroller that is able to process a limited amount of sensor data and make autonomous navigation decisions. Each robot has the ability to localize other robots using an RF-sonar combination for trilateration, and to communicate sensor readings using low-power, short-range radio transceivers. The distributed, decentralized nature of our design endows it with several desirable properties, including self-organization, self-repair, and efficient resource utilization (Spears *et al.*, 2004). These desirable properties are achieved despite sensor noise and latency, non-holonomic drives, limited velocity and acceleration capabilities, localization noise, and even in the presence of occasional hardware failures. Furthermore, the environment was surprisingly uncontrolled during the experiments – people walked into and out of the room in which the experiments were being conducted, windows leaked airflow into the room, and the robots unexpectedly shoved the emitter on occasion. Despite these hindrances, as well as exceptionally stringent metrics for success, our success rate is on par with that of more controlled experiments found in the literature. Spears *et al.* (2006) describe indoor laboratory experiments demonstrating the success of CPT (chemotaxis) on the real robots.

Research is currently in progress to identify and integrate anemometers with the robot platforms because in the near future we plan to implement anemotaxis and fluxotaxis on the robots as well. Anemotaxis is a strategy for traveling upwind. (Note that even though for simplicity of presentation we sometimes use gas-specific terms such as "upwind" or "airflow," the CPT algorithms apply to aquatic environments as well. In other words, anemotaxis vehicles travel upstream when in water.) Indoor and outdoor experiments will be performed to compare the three CPT strategies under a wide variety of real-world conditions.

4. CPT as an application of fluid dynamics

We have found that the foundation of effective CPT algorithms can be naturally and effectively built upon the fundamentals of fluid dynamics. This section shows the relationship between fluid physics and CPT. Note that the term "fluid" refers to both gaseous and liquid states of matter. The theoretical results presented in this paper are applicable to both gases and liquids. Therefore, our findings are applicable to the design of ground, marine, and aerial autonomous vehicles.

A physical description of a fluid medium consists of density, velocity, specific weight and gravity, viscosity, temperature, pressure, and so on. Here, due to their relevance and importance to the goals of CPT, we focus on the scalar density ρ, which has the units of mass per unit volume, and the vector velocity $\vec{V}$. Since our focus is on CPT, we use the term "density" and its symbol ρ to denote the concentration of the chemical being traced, rather than the density of the carrier fluid. For simplicity of presentation, unless stated otherwise, this paper assumes that the ambient fluids are compressible because in reality all materials are to some extent compressible. Furthermore, relative to a plume-tracing vehicle in motion, the distinction between the velocities of the fluid-borne chemical and that of the chemical-bearing fluid is often well below the sensor resolution threshold. In our model of the chemical plume, we assume that the chemical and ambient fluid velocities are identical.

Consider an infinitesimal fluid element moving within the plume. The velocity of this element can be represented as the vector $\vec{V} = u\hat{i} + v\hat{j} + w\hat{k}$, where $\hat{i}$, $\hat{j}$, and $\hat{k}$ are the

orthonormal basis vectors along the x-, y-, and z-axes in $\mathbb{R}^3$, i.e. $\vec{V} = < u, v, w >$. Note that $u = u(x, y, z, t)$, $v = v(x, y, z, t)$, $w = w(x, y, z, t)$, and t denotes time. Since chemical flow is typically unsteady, $\partial \vec{V} / \partial t \neq 0$ in the general case. The density and other characteristics of the fluid, i.e. the "flow-field variables," are likewise functions of space and time.

Sensor measurements of these flow-field variables are collected by the robots. In particular, it is assumed that a swarm of mobile vehicles equipped with chemical and fluid flow sensors can sample the plume to determine the current state of the contamination, and then use their on-board computers to collaboratively decide the most effective direction in which to move next. The calculations and collaborative decisions can be effectively based upon the physical rules that govern the flow of fluids, as we will demonstrate shortly.

The three governing equations of computational fluid dynamics for modeling fluid flow are the continuity, momentum, and energy equations. The continuity equation ensures that mass is conserved, the momentum equation captures Newton's second law $\vec{F} = m\vec{a}$, and the energy equation enforces energy conservation. Each equation can be expressed in at least four different forms, depending on one's perspective with respect to the fluid:

(1) a finite control volume fixed in space with the fluid moving through it;

(2) a finite control volume moving with the fluid;

(3) an infinitesimal fluid element fixed in space with the fluid moving through it; or

(4) an infinitesimal fluid element moving with the fluid (Anderson, 1995).

Selecting the third perspective, which is the most intuitive formulation for CPT, we write the conservation of mass equation as:

$$\frac{\partial \rho}{\partial t} = -\vec{\nabla} \cdot (\rho \vec{V}). \tag{1}$$

One can see that this equation expresses the physical fact that the time rate of decrease of mass inside the differential element (left-hand side) must equal the net mass flux flow (right-hand side) out of the element. Since we focus on numerical, rather than analytical, solutions to the fluid equations for use by computationally restricted robots, we retain the fundamental forms of these equations, rather than adopt a full Navier-Stokes formulation.

To model fluid flow, the equations are solved using numerical methods. In the case of perspective (3), the method of finite differences is often employed. For the CPT problem, a distinction must be made between the forward solution and the inverse solution. The former consists of a simulation that models the chemical flow, and the latter requires tracing within the fluid to find the source emitter, i.e. the latter is CPT. For the purposes of developing and testing new CPT approaches in simulation, the forward solution is required. Once the CPT algorithm has been refined in fluid-dynamic simulations, then the forward solution is no longer necessary, since the robots are tested with real laboratory-scale chemical plumes (Figure 2). However, note that the inverse solution (i.e. the CPT algorithm) is required for both the simulation and laboratory testing.

A numerical model of the forward solution for a given fluid provides data at discrete points in space, i.e. at "grid points." Abstractly, one can see a relationship between grid

points and robots in a swarm. This abstract correspondence motivates our view of a swarm of robots as a distributed computer that jointly senses the flow-field variables, shares them with their immediate neighbors, and then decides in which direction to move in order to locate the chemical emitter, thereby performing embodied computation. This relationship will be exploited, as relevant, in the remainder of the paper.

Consider one very important variable to measure in the context of CPT – the divergence of the velocity of the fluid (and therefore of the chemical). Note that as a fluid flows, any element within that fluid has invariant mass, but its volume can change. The divergence of the velocity, $\vec{\nabla} \cdot \vec{V}$, is the time rate of change of the volume of a moving fluid element, per unit volume. The divergence is the rate at which a fluid expands, or diverges, from an infinitesimally small region. A local vector field with positive divergence expands, and is called a source; a vector field with negative divergence contracts, and is called a sink. The magnitude of the divergence vector, $|\vec{\nabla} \cdot \vec{V}|$, is a measure of the rate of expansion or contraction. A vector field with zero divergence, when measured over the entire velocity vector field, is called solenoidal. The definition of solenoidal implies fluid incompressibility. This can be understood by considering the $\vec{\nabla} \cdot \vec{V} = 0$ constraint on the conservation of mass (1), where $(\partial/\partial t + \vec{V} \cdot \vec{\nabla})\rho = 0$, so that the total time derivative is zero, $\mathrm{d}\rho/\mathrm{d}t = 0$, which indicates that the fluid density is constant when viewed along the trajectory of a moving fluid element. Hence, compressible fluid flow with sources and sinks is not only realistic, but is also desirable – because the sources and sinks can facilitate local navigation in the absence of global landmarks (Decuyper and Keymeulen, 1991). Most importantly, a mathematical source (in terms of divergence) that remains a source over time (i.e. it is not transient) indicates a chemical emitter, whose identification is the key to successful CPT. This will be discussed in-depth below, because it is the reason why CPT is naturally considered as an application of fluid dynamics.

First, however, we need to broaden the notion of divergence to include the chemical density, in addition to the wind velocity. The key notion we seek is that of chemical mass flux, or mass flux for short. The mass flux is the product of the (chemical) density and the velocity, i.e. $\rho\vec{V}$. Informally, this is "the stuff spewing out of the chemical emitter." An astute reader will now infer that what would be of great interest for CPT is the divergence of the mass flux, which is, in mathematical notation in 3D:

$$\vec{\nabla} \cdot (\rho\vec{V}) = \vec{V} \cdot \vec{\nabla}\rho + \rho\vec{\nabla} \cdot \vec{V} = u\frac{\partial \rho}{\partial x} + \rho\frac{\partial u}{\partial x} + v\frac{\partial \rho}{\partial y} + \rho\frac{\partial v}{\partial y} + w\frac{\partial \rho}{\partial z} + \rho\frac{\partial w}{\partial z}. \qquad (2)$$

The divergence of the mass flux is the time rate of change of mass per unit volume lost at any spatial position. If this divergence is positive, it indicates a source of mass flux; if negative, it indicates a sink of mass flux. Sustained (over a period of time) positive divergence of chemical mass flux implies chemical spewing outward from a chemical source emitter, as opposed to a transient source.

To conclude this section, consider two modes of fluid transport: diffusion and advection. Diffusion dominates in an indoor setting where the air is relatively stagnant, e.g. when the windows are closed and there is little disturbance in the room. It also dominates at smaller spatial scales, e.g. insects tend to be more sensitive to diffusive effects than larger animals (Crimaldi *et al.*, 2002). In the absence of air flow, diffusion of chemical mass away from a chemical source will result in a Gaussian density profile. Advection is a more macroscopic phenomenon than diffusion. Both are at work to

disperse the chemical away from an emitter, but they have different effects on the chemical density profile that makes up a chemical plume. In this paper, we examine effects of both diffusion and advection processes on CPT.

5. Divergence-based signature of a chemical emitter

A key fluid dynamics concept that enables identification of a source of chemical flow is the divergence of the chemical mass flux. Recall from equation (2) that the divergence of mass flux (DMF), $\vec{\nabla} \cdot (\rho \vec{V})$, is the time rate of change of mass per unit volume. The mass flux combines chemical density and the velocity of chemical flow into one term. A positive divergence of this quantity implies a mathematical source. In a more practical sense, a source may be either transient or sustained. To differentiate a true source associated with the chemical emitter from a transient source that is merely a puff of chemical created by turbulence or other chemical distribution phenomena, we require sustained positive divergence of the mass flux in order to identify a clear emitter signal. In summary, the physical signature of a chemical emitter is $\vec{\nabla} \cdot (\rho \vec{V}) > 0$ for a period of τ experimentally determined time units.

To detect this signature and to perform CPT, the only additional sensors (beyond standard sensors, like IR for obstacle avoidance) required for robots are chemical sensors to detect the density ρ and anemometer sensors to detect the velocity $\vec{V}$. Sensors for discriminating which chemical is present are not discussed here. We assume they have already been applied to determine the chemical class. This paper focuses instead on robots that need to determine the value of ρ, i.e. chemical density rather than type, for a known contaminant.

6. The inverse CPT problem from a fluids perspective

Computational fluid dynamics is adept at modeling the forward solution of real fluid flows for the purposes of theoretical analysis and engineering design. CPT, on the other hand, seeks a solution to the inverse problem. Assuming that at least one source emitter is present in the environment (for simplicity, herein we assume a single emitter, though generalization to multiple emitters is straightforward)[1], finding that source requires environmental information. In an ideal situation, the CPT robots would have access to global knowledge, such as geo-referenced coordinates of the emitter, which can then be used to navigate directly to the source. But in most realistic situations, a local navigation strategy is the best that the robots can do, and that is what is assumed in this paper.

The CPT task consists of three subtasks:

(1) detecting the toxic chemical;

(2) tracing the chemical to its source emitter; and

(3) identifying the source emitter.

We consider each of these subtasks, in turn[2].

6.1 The chemical detection subtask: exploration

At the initiation of CPT, it is frequently assumed that the robots are not in contact with the chemical. Therefore, their first job is to locate the chemical plume. Under the assumption that they have no background knowledge about where the emitter is,

the best they can do is to engage in an undirected exploratory search. In the literature, this search typically takes the form of a zigzag or spiral pattern (Hayes *et al.*, 2001; Russell *et al.*, 2000). It is called casting because the robots are essentially casting about the environment, looking for a detectable (i.e. above sensor threshold) amount of the chemical. Casting strategies are often motivated by observation of biological creatures, such as insects, birds, or crustaceans (Grasso and Atema, 2002).

The chemical detection subtask is not described further in this paper. Although it is a very important part of CPT, the focus of this paper is on CPT from a fluid dynamics perspective, and casting strategies are not fluids-based.

6.2 The plume tracing subtask: following local signals

The objective of the plume tracing subtask is to follow clues, which are typically local if robots are used, to increase one's proximity to the source emitter. We have identified five major classes of algorithms for this subtask:

(1) *Heuristic strategies*.

(2) *Chemotaxis strategies*. Usually, these algorithms follow chemical gradients.

(3) *Infotaxis strategies*. There is a wide variety of algorithms in this class. A popular version follows the frequency of chemical puffs (i.e. filamentary structures) which typically increases in the vicinity of the emitter.

(4) *Anemotaxis strategies*. These algorithms move the robots upwind.

(5) *Hybrid strategies*. The more modern tracing approaches combine both chemical and upwind strategies.

More details, along with references, will be provided in Section 8. First, however, we elaborate a bit on the motivation for chemotaxis, infotaxis, and anemotaxis.

The chemotaxis strategy of chemical gradient following is motivated primarily by the chemical diffusion process that is responsible for some of the dispersion of the chemical from the emitter. Recall that diffusion creates a Gaussian chemical profile, i.e. the chemical density is highest near the emitter and it is reduced exponentially as distance from the emitter increases. In particular, the diffusive Gaussian chemical density profile can be modeled using the density diffusion equation, which is also simply called the diffusion equation. In its simplest 1D form, this equation is:

$$\frac{\partial \rho}{\partial t} = D \frac{\partial^2 \rho}{\partial x^2}.$$

This equation describes the time evolution of the density of chemical particles, $\rho(x, t)$, as a function of position, x, and time, t. It is particularly useful to consider a temporal initial condition ($t = 0$) when all the chemical is contained at the origin ($x = 0$), represented as the Dirac delta function particle source, i.e. $\rho(x, 0) = N\delta(x)$, where N is the total number of chemical filaments. $\delta(x)$ is the delta function, which has the value 0 everywhere except at $x = 0$, where its value is infinitely large and the total spatial integral is one. D is the diffusion coefficient, which is proportional to the square of the mean free path, λ_{mfp}, divided by the collision time, τ_{coll}, of particles in the fluid, that is, $D \propto \lambda_{\text{mfp}}^2 / \tau_{\text{coll}}$. The mean free path is the average distance traveled by particles in the fluid before they collide. The collision time is the average time between such collisions.

The solution to the 1D density diffusion equation, with a delta function initial condition, is the Gaussian evolution function:

$$\rho(x,t) = \frac{N}{\sqrt{4\pi Dt}}\, e^{-(x-x_o)^2/4Dt},$$

which shows exponential decay as a function of the distance x from the emitter. The emitter is assumed to be located at x_o (Figure 3 in Section 9.1). For a fixed time t, the solution fits a Gaussian function profile. Environments dominated by diffusion will be discussed further in Section 10.

Although chemical density following was one of the earliest chemical tracing strategies, others have also been explored. Consider infotaxis. It has been noted in outdoor settings (where advection is a stronger dispersive force than diffusion) that chemical density does not smoothly drop off with distance from the emitter. Instead, in outdoor scenarios it is common to have "puffs" of high-concentration chemical interspersed with little or no detectable chemical (Grasso, 2001; Iacono and Reynolds, 2008). This is due to the bulk turbulent diffusivity of the air. In this case, it appears that a relatively reliable signal to trace is the chemical intermittency, rather than the chemical density. The chemical intermittency is defined as the percentage of time that a sensor can detect an above-threshold chemical density. It is expected that intermittency increases proportionally with proximity to the source emitter. Intermittency-driven strategies form an important subclass of infotaxis.

Advection is a physical fluid process that motivates the anemotaxis strategy of traveling upwind. Since advection transports the chemical downwind, the robots expect to get closer to the origin of the chemical emission by moving upwind.

Considering the practical aspects of robotic CPT, observe that anemotaxis does not require any sharing of information between robots. Therefore, it is a strategy that could be employed by either a single robot or a collective of independent robots. For an implementation of chemotaxis, one robot with multiple spatially separated chemical sensors would suffice, but a much more efficient and effective approach uses a swarm of tens or hundreds of robots (Zarzhitsky and Spears, 2005). Section 8 describes several variations of the five major classes of strategies outlined in this section.

6.3 The emitter identification subtask: signature recognition

A chemical emitter leaves a physical signature by which it may be identified, whether or not CPT algorithms take advantage of this signature. This signature, given by equation (2), is the DMF, $\vec{\nabla}\cdot(\rho\vec{V})$, as has been described above.

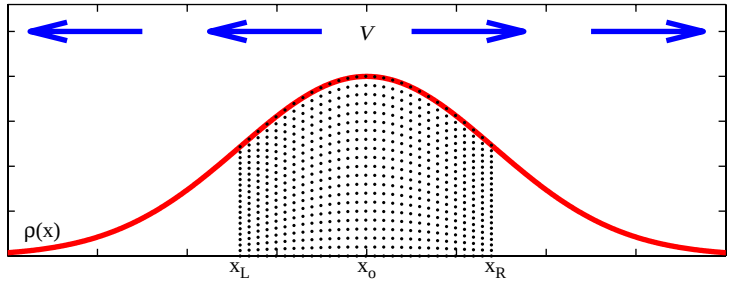

Note: The shaded area $x_L \le x \le x_R$ indicates the region where plume tracing is carried out by the fluxotaxis robots, and x_o marks the location of the chemical emitter

Figure 3.
The Gaussian chemical density distribution and the constant radial outflow velocity profile used in the constant flow speed theorem

Of course, the type of emitter will also have an effect on the character of its signature. There are continuous emitters, as well as single-puff emitters, and regularly and irregularly pulsing emitters. In this paper, we will address single-puff and continuous emitters – because real-world CPT problems typically contain a source that is a combination of these basic emission modes. Nevertheless, note that these are just examples; our theory is general and is not restricted to these two types of emitters. The size, shape, orientation, and height of an emitter can also affect the plume characteristics. The theoretical results in this paper are deliberately general, but could be parametrized as needed to model particular emitter geometries and configurations.

7. The major strategies for the emitter identification subtask

7.1 Heuristic strategies

One of the most popular heuristic approaches to emitter identification is a comparison of the chemical density at different heights (Cowen, 2002). The work of Grasso and Atema (2002) assumes that the emitter is low to the ground; they declare discovery of the emitter when the bottom sensor on their robot reports chemical detection for 90 percent of the sampling period. This percentage was empirically determined with experiments. On the other hand, Buscemi *et al.* (1994) determine that the emitter has been reached when the chemical sensors become saturated. Note that these heuristic approaches are prone to a high rate of false alarms (i.e. erroneously labeling sustained regions of high chemical concentration, such as areas next to walls or obstacles, as the source).

7.2 Machine learning strategies

As an alternative to heuristic techniques, the emitter identification subtask can be performed using machine learning approaches. For example, Lilienthal *et al.* (2004) treat this as a classification problem and use neural networks and support vector machines to learn to identify a chemical source. The advantage of this approach is that after training, they were able to achieve approximately 85 percent of the maximum hit rate. The disadvantage is that it requires a training phase, which makes the perhaps unrealistic assumption that the emitter type on which training occurs will be similar to the emitter type to be encountered during the actual CPT mission.

The approach adopted by Weissburg *et al.* (2002) could also be considered a machine learning approach. Their system is trained on characteristic emitter and non-emitter chemical patterns; it is then able to detect the emitter more accurately.

7.3 Emitter identification with fluxotaxis

Recall from Section 5 that the physical signature of a chemical emitter is $\vec{\nabla} \cdot (\rho \vec{V}) > 0$ for a period of τ (experimentally determined) time units. How could this signature be identified by a team of robots trying to solve the emitter identification problem? Fluxotaxis employs the divergence theorem of vector calculus (Hughes-Hallett *et al.*, 1998):

$$\int_{W_o} \vec{\nabla} \cdot (\rho \vec{V}) \mathrm{d}W = \oint_{A_o} (\rho \vec{V}) \cdot \hat{n} \mathrm{d}A. \tag{3}$$

Simply put, the left side of this equation is a volume integral of the DMF, while the right side of the equation is the surrounding surface area integral of the mass flux in the outward normal direction. This equation, where W_o is the control volume, A_o is the bounding surface of the volume, and $\hat{n}$ is a unit vector pointed as an outward normal to

the surface of the volume, allows us to formally define the intuitive notion that a control volume containing a source will have a positive mass flux divergence, whereas a control volume containing a sink will have a negative mass flux divergence. This result can be used as the basic criterion for theoretically identifying a chemical emitter, which is a true mathematical source under this definition. In particular, equation (3) shows that if the robots encircle a suspected emitter, and the total mass flux exiting the circle of robots consistently exceeds some small, empirically determined threshold for a given amount of time, then the robots are known to surround a true chemical emitter. Our fluxotaxis algorithm uses this theoretical criterion. To the best of our knowledge, previous criteria for emitter identification are purely heuristic (Cowen, 2002). Fluxotaxis is the first with a firm theoretical basis.

8. The major strategies for the plume tracing subtask

Some of the earliest solutions of the CPT problem adopted a time-averaged approach for identifying the chemical; however, more recent research indicates that mean statistics converge too slowly (Farrell *et al.*, 2002; Liao and Cowen, 2002).

Early robotic experiments based on solutions of fluid dynamic problems are reported by Decuyper and Keymeulen (1991). In their approach, a simplified model of fluid flow formed the basis of a simulated robot's navigation strategy. This method, refined further by Keymeulen and Decuyper (1994a, b), was inspired by the fact that continuous fluid flow can be used for iterative optimization of the local-to-global route finding task, since the pressure fields responsible for stable optimal flow paths are void of local minima. A successful development of this strategy also relied on the concepts of a fluid source and sink, which they used to represent the robot's position. However, the task addressed by Decuyper and Keymeulen (1991) consisted of navigation within an obstacle-filled room, and since it did not involve emitter localization, their results are not immediately applicable to the CPT problem.

8.1 Simple heuristic approaches

Two of the most popular heuristic approaches are following the plume centerline and following the edge of the plume (Cowen, 2002). In fact, these two heuristic approaches may be considered complementary. For example, by following the edge of the plume, one may more rapidly identify and move into the plume centerline.

8.2 Chemotaxis

Chemotaxis is the best understood and most widely applied biomimetic CPT approach. It consists of tracing the chemical, typically by following a local gradient of the chemical concentration within a plume (Krishnanand and Ghose, 2006; Marques *et al.*, 2006). Some of the earliest research on chemotaxis was performed by Sandini *et al.* (1993) at the University of Genoa. Among the most extensive applications of chemotaxis are those of Lilienthal *et al.*. In some of their work they show chemotaxis success in an uncontrolled indoor environment (Lilienthal and Duckett, 2003). They have also explored chemotaxis in ventilated corridors with weak chemical sources (Lilienthal *et al.*, 2001).

Feddema *et al.* (2003) apply a control theoretic approach to chemotaxis that allows for stability and convergence proofs with optimality guarantees. On the other hand, it makes a strong assumption that the chemical plume density profile can be closely approximated by a quadratic surface.

While chemotaxis is very simple to perform, it frequently leads to locations of high concentration in the plume that are not the real source, e.g. a corner of a room (Song and Chen, 2006). Cui *et al.* (2004) investigated an approach to solving this problem by using a swarm majority vote, along with a communication and routing protocol for distributing information to all members of a robotic collective. However, that is a strong requirement, and Cui *et al.* also make an even stronger assumption that each robot in the collective has a map of the environment.

In addition to the local maxima problem of chemotaxis, we show in Section 10 that a chemotaxis search strategy can fail near the emitter's location, due to the fact that for a typical, time-varying Gaussian distribution profile, the chemical density gradient goes to zero near the distribution's peak.

8.3 Infotaxis

There are different varieties of infotaxis. One, which we call intermittency infotaxis, is similar to chemotaxis because it considers only the chemical density and disregards the ambient wind velocity. On the other hand, it is different from chemotaxis because it measures the chemical intermittency (i.e. the "puff" frequency), rather than density gradients, to determine source proximity. Liao and Cowen (2002) demonstrate the advantages of intermittency infotaxis over alternative approaches. Related to intermittency is the notion of coherency. Kikas *et al.* (2001) measure coherency using a correlation analysis of temporal concentration fluctuations across a sensor array.

Justus *et al.* (2002) explore not only intermittency, but also the burst length, where "burst" is another name for a puff or filament. The burst length is found to be a finer-scale measure of source proximity than intermittency. In particular, frequent short bursts and few long bursts both increase the intermittency. The intermittency is high near a chemical source, but it may also be high in other regions of the plume. Burst length enables discrimination of such cases.

A second type of infotaxis focuses on information theory and entropy. Vergassola *et al.* (2007) have developed a strategy that locally maximizes the expected rate of information gain, in an information theoretic sense. Actually, their approach is a Bayesian and information theoretic variant of intermittency infotaxis. Another probability-driven approach to infotaxis is the maximum likelihood approach of Balkovsky and Shraiman (2002). They focus on carrier wind velocity, and maximize the conditional probability $P(source|sequenceofobservations)$. Jeremić and Nehorai (1998) also use a maximum likelihood approach, but they use it for tuning parameters. These parameter values are learned for the purpose of instantiating a sensor model that takes into account chemical measurements, signal bias and Gaussian noise.

Parunak and Brueckner (2001) conduct analysis of the self-organization property in multi-robot systems from the standpoint of entropy and the second law of thermodynamics. They develop an analogy between entropy and information disorder, and show how understanding and control of system entropy can be used to organize a multi-robot system. They illustrate the idea by solving a robot coordination problem with the use of simulated randomly diffusing pheromones.

8.4 Anemotaxis

Another common approach, usually called "anemotaxis" but sometimes alternatively called odor gated rheotaxis, has been proposed for the CPT task. An anemotaxis-driven

robot focuses on the advection portion of the flow. It measures the direction of the fluid's velocity and navigates "upstream" within the plume. Hayes *et al.* (2001) have done seminal work in this area. Grasso and Atema (2002) combine anemotaxis with casting, so that the robots alternate moving upstream and cross-stream in the wind. The simulation results of Iacono and Reynolds (2008) show that the effectiveness of anemotaxis improves with increased wind speed. More complex wind-driven strategies may be found in Kazadi *et al.* (2000). They have successfully explored a form of anemotaxis that combines wind velocity information with passive resistive polymer sensors and strategic robot placement which enables effective anemotaxis. Ishida *et al.* (2006) gain performance improvements by coupling anemotaxis with vision capabilities.

One of the biggest advantages of anemotaxis is that it can be performed successfully with either a single robot, or with a group of independent robots. The same also holds for infotaxis. However, Grasso and Atema (2002) have found that two spatially separated wind sensors outperform a single wind sensor. Although anemotaxis can be a very effective strategy for some problems, its limitation is that it can lead to a wind source that is not the chemical emitter.

8.5 Heuristic hybrids

Recently, there has been an increasing trend toward the development of hybrid CPT strategies that combine chemotaxis with some form of anemotaxis. For example, Ishida *et al.* (2001) use a simple hybrid strategy that consists of applying anemotaxis when the chemical density ρ is high, and otherwise applying chemotaxis.

Russell *et al.* (2000) employ a more complex hybrid approach to CPT for finding the source of a chemical emission in a maze. Their single CPT robot employs a heuristic algorithm for traveling upwind, following the chemical and avoiding obstacles.

More unusual hybrids are the approaches adopted by Lytridis *et al.* (2006), Marques *et al.* (2006). They combine chemotaxis with a biased random walk, and discover that this hybrid outperforms its components.

8.6 The fluxotaxis hybrid

Fluxotaxis is the most fluid-based of all the strategies, and therefore the remainder of this section focuses on fluxotaxis. Recall the definition of the divergence of the mass flux from equation (2). This divergence is the hallmark signature of a chemical source emitter. A logical conclusion is that following its gradient will take the robots closer to the source. That is the essence of our fluxotaxis strategy, i.e. to follow the gradient of the divergence of mass flux, abbreviated *GDMF*, where the gradient is the direction of steepest increase. The mathematical formula for the *GDMF* in 3D is:

$$\vec{\nabla}\left[\vec{\nabla}\cdot(\rho\vec{V})\right] = \vec{\nabla}\left(u\frac{\partial\rho}{\partial x} + \rho\frac{\partial u}{\partial x} + v\frac{\partial\rho}{\partial y} + \rho\frac{\partial v}{\partial y} + w\frac{\partial\rho}{\partial z} + \rho\frac{\partial w}{\partial z}\right).$$

$\overrightarrow{GDMF}$ combines information about both velocity and chemical density, in a manner motivated by the theory of fluid dynamics. Mathematically, the DMF can be subdivided into two terms:

$$\vec{\nabla}\cdot(\rho\vec{V}) = \rho(\vec{\nabla}\cdot\vec{V}) + \vec{V}\cdot(\vec{\nabla}\rho). \tag{4}$$

Therefore, the $\overrightarrow{GDMF}$ is a gradient $\vec{\nabla}$ of the sum of two terms:

(1) $\rho(\vec{\nabla}\cdot\vec{V})$, which is the density times the divergence of the velocity field; and

(2) $\vec{V}\cdot(\vec{\nabla}\rho)$, which is the flow velocity field in the direction of the density gradient.

When the chemical flow is divergent, the first term takes precedence in directing the robots, and it is analogous to anemotaxis. When the fluid velocity is constant, or the flow is stagnant, then $\vec{\nabla}\cdot\vec{V}$ is zero and the second term is the only relevant term. This second term is the flow velocity in the direction of the density gradient, and is analogous to chemotaxis. What is important here is that $\overrightarrow{GDMF}$ theoretically combines chemotaxis and anemotaxis, and the relevant term is activated automatically based on the environment. No deliberate strategy selection or switching mechanism is required – $\overrightarrow{GDMF}$ automatically becomes whatever it is supposed to be in an environmentally driven manner.

Next, consider the relationship between the emitter signature and $\overrightarrow{GDMF}$ as a navigation strategy to locate the source emitter. In particular, suppose we apply $\overrightarrow{GDMF}$ as an enclosure fluxotaxis algorithm. In other words, assume that there is a robot lattice consisting of hundreds of robots surrounding the emitter, thus creating a full sensing grid. Also, assume that the chemical source within the control volume, W_o, contains a total amount of chemical, Q, at time t so that when the source emitter ejects the chemical (in a chemical puff or continuous fashion), the time rate of change of the total chemical is negative, $\mathrm{d}Q/\mathrm{d}t < 0$. Then we have a formal guarantee regarding $\overrightarrow{GDMF}$ as an enclosure navigation strategy. In particular, a key diagnostic of the chemical plume emitter location is when we have $\oint_{A_o}(\rho\vec{V})\cdot\hat{n}\mathrm{d}A > 0$, where A_o is the surface area of the volumetric region W_o that the robots are known to surround. If $Q = \int_{W_o}\rho\mathrm{d}W$ is the total chemical in the control volume W_o, then from the conservation of mass equation (1), integrated over the control volume, while using the divergence theorem equation (3), the time rate of change of the total chemical:

$$\frac{\mathrm{d}Q}{\mathrm{d}t} = -\oint_{A_o}(\rho\vec{V})\cdot\hat{n}\mathrm{d}A,$$

is indeed found to be negative, i.e. $\mathrm{d}Q/\mathrm{d}t < 0$, which is indicative of a source. In other words, a robot lattice surrounding the emitter can locate it by going opposite to the direction of flux outflow. In this case, we have a theoretical guarantee that the enclosure $\overrightarrow{GDMF}$ strategy will allow the robot lattice to home in on the source emitter. In order to backtrack where the source of the chemical plume originated, it is reasonable to look in the direction of the $\overrightarrow{GDMF}$ vector.

At this point, it is important to merge the global and local perspectives being maintained simultaneously above and in the remainder of this paper. From the global perspective, the reality (from both a fluids and a fluxotaxis perspective) is that we are dealing with control volumes for which we calculate the surface integral of the mass flux. Under typical real-world transient fluid conditions, within volumes near the chemical emitter the flux inflow frequently does not equal the flux outflow. It takes time for the chemical mass to move away from the emitter prior to reaching steady state (at which point there is no CPT signal for the robots to track). During the transient period of fluid emission, there are numerous volumetric regions that mimic "sources" and "sinks," sometimes far away from the emitter, where the chemical accumulates temporarily. In other words, the plume consists of a highly dynamic flow

of chemical-laden fluid, which is moving unsteadily through the region around the emitter. Typically, this implies that in a significantly large region containing a chemical source, the surface area flux integral will be nonzero, thereby providing a gradient signal for the robots to follow. By following the regions that behave like sources, fluxotaxis-driven robots quickly, and easily find the true source emitter. In numerous computational fluid dynamics simulations we have found that our gradient-following fluxotaxis algorithm quickly and reliably locates the source emitter using these ambient chemical flux signals. Even in complex, simulated urban environments filled with many large buildings and narrow passageways, we have shown that fluxotaxis demonstrates rapid and reliable success, whereas the leading alternative approaches typically flounder and fail (Agassounon *et al.*, 2009).

The analogous local perspective is that of the $\overrightarrow{GDMF}$ defined above. The $\overrightarrow{GDMF}$ is an idealized, local concept that we use as a visualization tool to illustrate what fluxotaxis does. With $\overrightarrow{GDMF}$ we assume an infinitesimal analysis based on differential calculus that reasonably approximates the global volumetric perspective. Although in reality we have a finite grid (both in simulation and with the robots), we can compute an approximate solution by assuming that the $\overrightarrow{GDMF}$ can be calculated at every point in space. Mathematically, we are stating that $\vec{\nabla}\left[\oint_{A_o}(\rho\vec{V})\cdot\hat{n}\mathrm{d}A\right]$ becomes equivalent to $\vec{\nabla}\left[\vec{\nabla}\cdot(\rho\vec{V})\right]$ in the limit as we move from discrete grids to continuous space. The latter is the $\overrightarrow{GDMF}$ formula. The advantage of adopting the local $\overrightarrow{GDMF}$ perspective is that it simplifies our analysis without loss of correctness. Using $\overrightarrow{GDMF}$ we can clearly illustrate why fluxotaxis-guided robots move in the correct direction toward the source emitter.

To conclude this section, we briefly mention some practical considerations for applying fluxotaxis and other CPT methods. Any CPT strategy will only be valid in regions where the value of ρ and its gradients (if computed by the strategy) are above some minimally significant threshold. Below this minimum density (e.g. too far from the source emitter), all CPT algorithms will be poor predictors of the source location, and the robots should resort to casting. Fluxotaxis navigates best in the presence of sources and sinks, which are common in realistic fluid flows, and it performs well in all other CPT fluid conditions in which we have tested it. It is robust in its avoidance of sinks, and exceptionally precise in homing in on sources. Interestingly, we have experimentally determined that a local, neighbor-oriented, first-derivative approximation of $\overrightarrow{GDMF}$ works best for swarms of robots navigating in regions far away from sources and sinks, with above-threshold levels of chemical. In such regions, the density may not be sufficiently far above the threshold for a second derivative to be useful, and therefore a first-derivative approximation is more suitable. With this version of fluxotaxis, every robot calculates the component of mass flux, $\rho\vec{V}$, in the direction from each locally sensed neighbor to itself, and it moves in the direction in which this scalar value is maximized. This swarm-oriented first-derivative $\overrightarrow{GDMF}$ approximation, which improves as the swarm size increases, is discussed at length by Zarzhitsky (2008). We will not discuss it further in this paper because it is an implementation issue. Instead, this paper focuses exclusively on the precise, second-derivative, original $\overrightarrow{GDMF}$ formulation. Finally, for solenoidal regions in which the fluid is incompressible and the chemical density is invariant along velocity trajectories, heuristic CPT approaches or casting must be used. When the chemical has

dissipated for a long time and a steady state has been reached, there will be no chemical derivative signal for any CPT strategy to follow.

9. A theoretical fluid-based analysis of fluxotaxis

In the previous section, we presented multiple strategies for the plume tracing subtask of CPT. Most of these strategies are motivated by biological CPT systems, and thus do not use fluid dynamics principles for an emitter signature identification. The one exception is the fluxotaxis strategy, which is derived fully from the fluid dynamics properties of chemical plumes. Since this paper is about CPT from a fluid dynamics perspective, this section explores the physical underpinnings of fluxotaxis, and we focus on the gradient of the divergence of the chemical mass flux – a fundamental component of the fluxotaxis algorithm.

In this section, we prove a sequence of theorems that elucidate the strengths of $\overrightarrow{GDMF}$ as a local guide for finding the chemical emitter[3]. The assumptions may, at first blush, appear to be restrictive; however, they form the basis for some of the most realistic flow situations to be encountered in environmental scenarios. Furthermore, there is no assumption about the robot lattice surrounding the emitter. All of these theorems assume a shared coordinate system between neighboring robots[4]. The section assumes a single coordinate axis (1D) for simplicity of presentation – so that the reader can gain the necessary insight and intuition into our approach. In other words, to explain why $\overrightarrow{GDMF}$ is effective, we begin with a 1D tutorial perspective when presenting our theoretical results. Furthermore, for improved clarity, we focus on static "snapshots," or moments in time. However, we have verified the correctness of all results for the dynamic 2D and 3D cases. Section 9.4 summarizes the 3D version of this analysis.

Let us begin with a 1D environment, where the robot lattice consists of a linear formation of spatially separated robots. Each robot can detect the chemical density ρ and the wind velocity $\vec{V}$ in its vicinity, and calculate the $\overrightarrow{GDMF}$ via an exchange of sensor data with its neighbors. The objective of this section is to present theorems regarding whether the $\overrightarrow{GDMF}$ vector calculated by the grid robots "points toward" (is predictive of) the source emitter, or "points away from" a sink, under three classical environmental conditions:

(1) a constant flow speed;

(2) a chemical flow source; and

(3) a chemical flow sink.

Fluid compressibility is assumed for the source and sink results, but not for the constant flow results.

The technique employed in these theorems and proofs is to select two points, P_{emt} and P_{far}, where two robots are stationed within the lattice, so that the former is closer to the emitter, and the latter is farther from the emitter. At each time step, the robots share their DMF values, calculate the $\overrightarrow{GDMF}$, and then move in the direction of the $\overrightarrow{GDMF}$ vector. If $\text{DMF}_{\text{emt}} > \text{DMF}_{\text{far}}$ then the $\overrightarrow{GDMF}$ will predict the correct direction – toward the chemical source. When dealing with sinks, on the other hand, two points P_{snk} and P_{far} are identified and in this case the correct prediction occurs when $\text{DMF}_{\text{snk}} < \text{DMF}_{\text{far}}$. In other words, the correct robot behavior is to avoid sinks.

The two robots under consideration could be any pair of neighbors within the lattice that share their DMF values to determine the new direction in which to navigate.

9.1 1D GDMF and constant flow speed

Constant flow speed theorem. Assume that the following conditions hold:

- The chemical plume has a Gaussian distribution $\rho(x) = \kappa e^{-(x-x_o)^2/\Delta x^2}$, centered at x_o, where Δx is the density profile scale length.
- The lattice is positioned between x_L and x_R, where x_L and x_R are solutions to (Figure 3):

$$\frac{\partial^2 \rho(x)}{\partial x^2} = 0,$$

  this implies that $\partial^2\rho(x)/\partial x^2 < 0$ in the region of interest. This is the region where CPT can be effective.
- $\vec{V}$ is constant in magnitude throughout the flow, except right at the emitter (x_o), and is an outward radial vector.

Without loss of generality, assume the existence of robots at points P_{emt} and P_{far}, such that P_{emt} is closer to the emitter than P_{far}. Then execution of one step of the fluxotaxis algorithm implies that the robot lattice moves closer to the emitter, or equivalently:

$$\left[u\frac{\partial \rho}{\partial x} + \rho\frac{\partial u}{\partial x}\right]_{\text{far}} < \left[u\frac{\partial \rho}{\partial x} + \rho\frac{\partial u}{\partial x}\right]_{\text{emt}}. \tag{5}$$

Proof. The problem is symmetric with respect to the emitter's location (x_o); thus, it is sufficient to prove the case where $x_L < P_{\text{far}} < P_{\text{emt}} < x_o$. Since $|\vec{V}|$ is constant, $\partial u/\partial x = 0$, and equation (5) simplifies to:

$$\left[u\frac{\partial \rho}{\partial x}\right]_{\text{far}} < \left[u\frac{\partial \rho}{\partial x}\right]_{\text{emt}}.$$

Since u is a negative constant, the inequality can be simplified to:

$$\left[\frac{\partial \rho}{\partial x}\right]_{\text{far}} > \left[\frac{\partial \rho}{\partial x}\right]_{\text{emt}}.$$

Grouping like terms gives:

$$0 > \left[\frac{\partial \rho}{\partial x}\right]_{\text{emt}} - \left[\frac{\partial \rho}{\partial x}\right]_{\text{far}}.$$

This is true because, by Assumption 2:

$$0 > \frac{\partial^2 \rho}{\partial x^2}.$$

Results of a software simulation for this theorem are shown in Figure 4. In the figure, light-colored areas denote large values, and dark-colored areas correspond to small values. The location of the chemical emitter is marked by the triangle symbol. The initial positions of two separate robot lattices are at the outer edges of the environment,

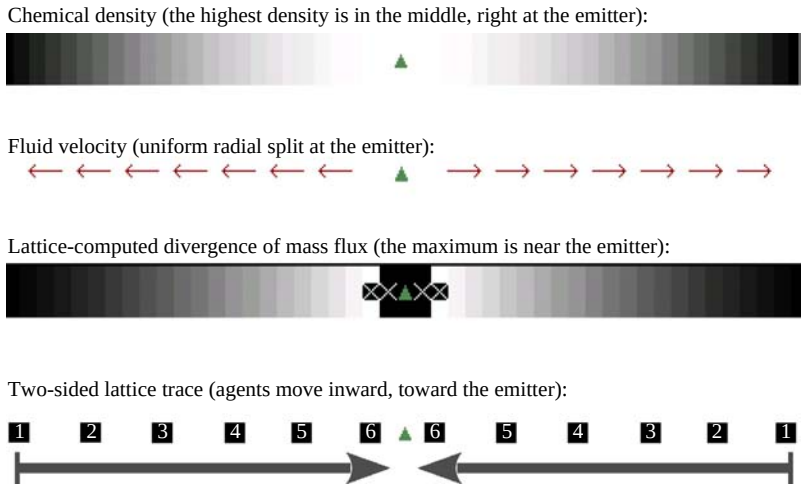

Notes: Individual robots are shown as black boxes with the white × in the middle, and the time trace of the two independent robot lattices is shown with boxed numbers indicating the location of the lattice at a given time step. The theorem holds for any initial lattice configuration, and fluxotaxis successfully locates the chemical emitter

Figure 4.
Simulation results for the constant flow speed theorem

to the left and right of the emitter. During execution of the fluxotaxis algorithm, each robot (shown as a black box with a white × in the middle) computes the DMF using equation (2), with the partial derivatives replaced by the second-order accurate central difference approximation (Zarzhitsky, 2008). This value is recorded by the simulator for analysis purposes, and is displayed along with the final robot positions in the screenshot. Observe that the resulting divergence "landscape" has a global peak that coincides with the location of the emitter, and does not have any local maxima that could trap or mislead the robots. There is a small gap in the computed divergence plot near the emitter because the robots had terminated their search upon reaching the emitter. Each simulated robot (the black box) corresponds to one of the reference points (P_{emt} or P_{far}) in the theorem's proof and, just as in the theorem, there are two robots per lattice. In this simulation, both robot lattices correctly moved toward the emitter in the center. In the proof of the constant flow speed theorem, we only considered the case where the lattice was to the left of the emitter; however, a similar proof can be given for the symmetric case, where the lattice starts out on the right side of the emitter, and the simulation in Figure 4 shows that the algorithm works regardless of the initial position of the robot lattice with respect to the position of the chemical source.

9.2 1D GDMF at a source

Source theorem. The fluxotaxis $\overrightarrow{GDMF}$ will advance the robot lattice toward a chemical source.

Proof. As before, assume a general Gaussian chemical plume distribution. Without loss of generality, assume the existence of two points P_{emt} and P_{far}, such that P_{emt} is closer to the source than P_{far} (Figure 5). Two cases result, based on the orientation of the lattice coordinate axis.

Case I. assumes that the direction of the lattice coordinate axis is opposite to the direction of the fluid flow, and thus:

(1) $\partial^2 u/\partial x^2 \geq 0$;

(2) $\partial u/\partial x > 0$; thus $0 \geq u_{\text{emt}} > u_{\text{far}}$;

(3) $\partial^2 \rho/\partial x^2 \leq 0$; and

(4) $\partial \rho/\partial x > 0$ and therefore $\rho_{\text{emt}} > \rho_{\text{far}}$.

We need to prove that the robot will move toward the source, or:

$$\left[u\frac{\partial \rho}{\partial x} + \rho\frac{\partial u}{\partial x}\right]_{\text{far}} < \left[u\frac{\partial \rho}{\partial x} + \rho\frac{\partial u}{\partial x}\right]_{\text{emt}}. \tag{6}$$

Assumptions 1 and 3 imply:

$$\left[\frac{\partial u}{\partial x}\right]_{\text{far}} \leq \left[\frac{\partial u}{\partial x}\right]_{\text{emt}} \quad \text{and} \quad \left[\frac{\partial \rho}{\partial x}\right]_{\text{far}} \geq \left[\frac{\partial \rho}{\partial x}\right]_{\text{emt}}.$$

Together, with Assumptions 2, 4, and algebraic rules, Case I holds. □

Case II. is with the lattice coordinate axis in the same direction as the fluid flow, so that both u_{emt} and u_{far} are non-negative (Figure 5), and the previous assumptions become:

(1) $\partial^2 u/\partial x^2 \leq 0$;

(2) $\partial u/\partial x > 0$; thus $0 \leq u_{\text{emt}} < u_{\text{far}}$;

(3) $\partial^2 \rho/\partial x^2 \leq 0$; and

(4) $\partial \rho/\partial x < 0$ and therefore $\rho_{\text{emt}} > \rho_{\text{far}}$.

The robot will turn around and move toward the source if equation (6) holds. From Assumption 1, we conclude:

$$\left[\frac{\partial u}{\partial x}\right]_{\text{far}} \leq \left[\frac{\partial u}{\partial x}\right]_{\text{emt}}.$$

Similarly, Assumption 3 yields:

$$\left[\frac{\partial \rho}{\partial x}\right]_{\text{far}} \leq \left[\frac{\partial \rho}{\partial x}\right]_{\text{emt}}.$$

Algebraic application of the remaining assumptions shows that equation (6) holds. □

Software simulation of this theorem's configuration for both cases is shown in Figure 6. As before, two *GDMF*-driven lattices (represented by black boxes marked with the white × symbol) begin at the outer edges of the simulated environment, and move in

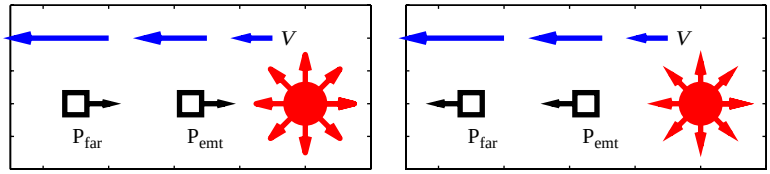

(a) Case I **(b)** Case II

Figure 5.
Local coordinate axis of the CPT vehicles and emitter location in the source theorem

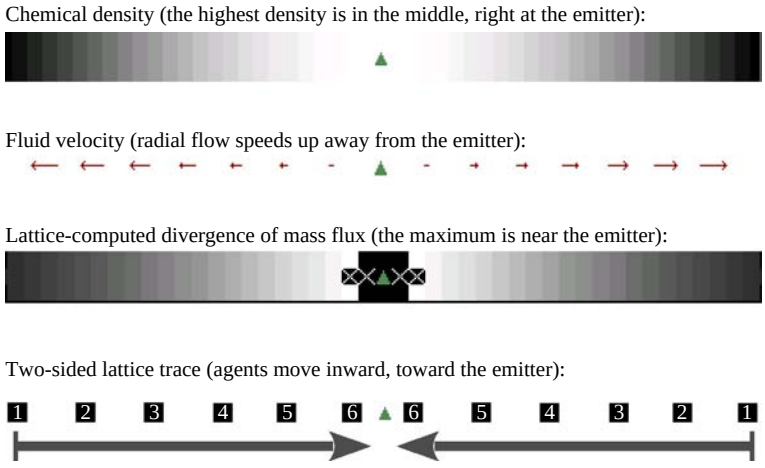

Notes: The time trace, denoted by the numbered boxes, shows the location of each of the two different robot lattices at sequential time steps in the simulation. Both lattices correctly converge on the true location of the chemical emitter

Figure 6.
Simulation of a GDMF-driven lattice (represented by black boxes) in the vicinity of a chemical source from the source theorem

toward the emitter, denoted by the triangle in the center. The direction of motion is determined by the gradient of the DMF, which is computed locally by each robot using a central difference approximation of the partial derivatives in equation (2), and as can been seen from the divergence plot, has the maximum value near the emitter's location. Similar to the previous simulation, the divergence value right at the emitter is not computed by the lattice, since the search terminates as soon as the emitter is found. Two fluxotaxis lattices are shown in the screenshot, and as expected, both of them successfully navigate toward the chemical source. As this figure illustrates, the initial position of a lattice with respect to the emitter does not affect the robots' ability to correctly localize the emitter. We should note here that in order to succeed, a robot lattice must discern between a transient versus a sustained source. An approach for distinguishing between transient and sustained sources is outlined in Section 5.

9.3 1D GDMF at a sink

Sink theorem. Fluxotaxis $\overrightarrow{GDMF}$-controlled robots will move away from a chemical sink (Figure 7).

Proof. As before, assume a general Gaussian chemical plume distribution. Without loss of generality, assume the existence of two points P_{snk} and P_{far}, such that P_{snk} is closer to the chemical sink than P_{far} (Figure 7). To prove that the robots will move away from the sink, we must show:

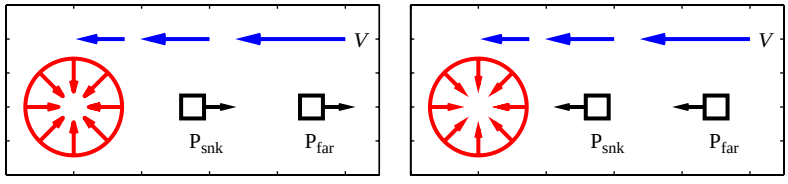

(a) Case I **(b)** Case II

Figure 7.
Location of the chemical sink and the two lattice orientations in the sink theorem

$$\left[u\frac{\partial \rho}{\partial x}+\rho\frac{\partial u}{\partial x}\right]_{\text{snk}} < \left[u\frac{\partial \rho}{\partial x}+\rho\frac{\partial u}{\partial x}\right]_{\text{far}}. \tag{7}$$

Two cases result, based on the orientation of the lattice coordinate axis.

Case I. Occurs when the lattice coordinate axis points in the opposite direction to the fluid flow, so that both u_{snk} and u_{far} are negative (Figure 7). For this case, the assumptions are:

(1) $\partial^2 u/\partial x^2 \geq 0$;

(2) $\partial u/\partial x < 0$; thus $0 \geq u_{\text{snk}} > u_{\text{far}}$;

(3) $\partial^2 \rho/\partial x^2 \leq 0$; and

(4) $\partial \rho/\partial x < 0$ and therefore $\rho_{\text{snk}} > \rho_{\text{far}}$.

The robot will continue moving away from the sink if equation (7) is true. From Assumption 1, we observe that:

$$\left[\frac{\partial u}{\partial x}\right]_{\text{snk}} \leq \left[\frac{\partial u}{\partial x}\right]_{\text{far}}.$$

Likewise, Assumption 3 implies:

$$\left[\frac{\partial \rho}{\partial x}\right]_{\text{snk}} \geq \left[\frac{\partial \rho}{\partial x}\right]_{\text{far}}.$$

The remaining assumptions with algebraic simplification prove that equation (7) is true. □

Case II. if the direction of fluid flow and the lattice coordinate axis are the same, then:

(1) $\partial^2 u/\partial x^2 \leq 0$;

(2) $\partial u/\partial x < 0$; thus $0 \leq u_{\text{snk}} < u_{\text{far}}$;

(3) $\partial^2 \rho/\partial x^2 \leq 0$; and

(4) $\partial \rho/\partial x > 0$ and therefore $\rho_{\text{snk}} > \rho_{\text{far}}$.

From Assumptions 1 and 3, we conclude that:

$$\left[\frac{\partial u}{\partial x}\right]_{\text{snk}} \leq \left[\frac{\partial u}{\partial x}\right]_{\text{far}} \quad \text{and} \quad \left[\frac{\partial \rho}{\partial x}\right]_{\text{snk}} \leq \left[\frac{\partial \rho}{\partial x}\right]_{\text{far}}.$$

Algebraic simplification using Assumptions 2 and 4 proves Case II. □

Simulation results for this theorem are shown in Figure 8. Confirming the theoretical results, the high-density chemical build-up in the center of the environment does not fool the fluxotaxis algorithm, which correctly avoids the local spike in the density by directing the robots (again represented by black boxes) to the outer edge of the tracing region, where as can be seen from the divergence plot, the maximum mass flux divergence occurs. The sink theorem proves that a *GDMF*-driven robot lattice will escape from a sink. However, a simple chemotaxis strategy is easily fooled by sinks, since by definition of a sink, $\partial \rho/\partial x > 0$ going into the sink. The fluxotaxis scheme is more robust in this case because it looks at the second-order partial derivative of ρ, and

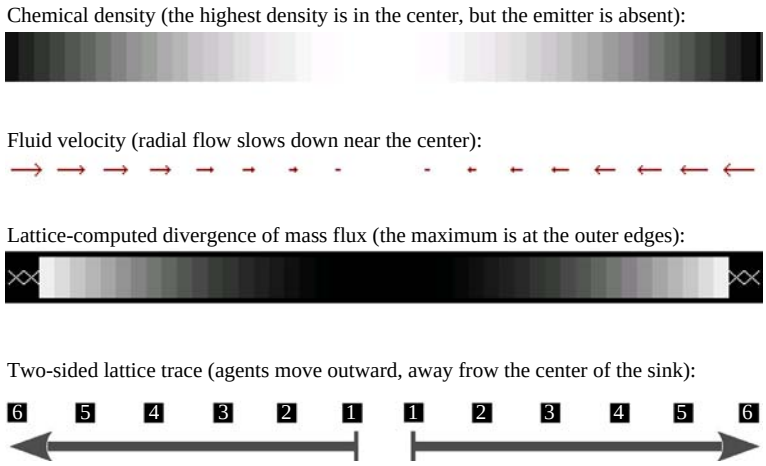

Notes: As stated in the proof and visualized in the last time-step diagram, the robust fluxotaxis method forces the robot lattice out of the sink, even if the lattice starts out directly in the center of the sink, where the chemical concentration is at a local maximum

Figure 8.
Simulated performance of the fluxotaxis algorithm within the chemical sink from the sink theorem

also takes the divergence of velocity into account. This simulation provides an example of how effectively the fluxotaxis technique merges the chemotaxis and anemotaxis CPT methods into a physically sound algorithm with valuable self-correcting properties.

9.4 3D GDMF extension of analysis

Although the *GDMF* source and sink analysis (given in Sections 9.1-9.3) has been presented in 1D, it should be noted that all results are also consistent in 3D for radially symmetric (isotropic) density and velocity profiles. Specifically, we assume a Gaussian chemical density profile, $\rho(\vec{r}) = \rho_o e^{-(r/a)^2}$, centered about a source at $\vec{r} = 0$, where ρ_o is the initial density at the source at time $t = 0$, a is the density scale length, and r is the magnitude of the vector $\vec{r}$, i.e. $r = |\vec{r}|$. We assume either a constant radial velocity profile, $\vec{V} = V_m\hat{r}$ (where $\hat{r}$ is a unit vector in the direction of $\vec{r}$, which points outward from the emitter), or an accelerated and asymptotic radially outward velocity profile, $\vec{V} = V_m[1 - e^{-(r/b)}]\hat{r}$, from the source point at $\vec{r} = 0$, where b is the velocity scale length and V_m is the maximum fluid speed attained at a large radius.

The first velocity profile scenario (Section 9.1), using a constant radial velocity profile in combination with a Gaussian radial density profile, is explored in detail in Section 10.1. Simply put, in this scenario, for "small r,"[5] where approximate asymptotic scalings are $\rho \sim \rho_o[1 - (r/a)^2]$, $\partial\rho/\partial r \sim -2\rho_o r/a^2$, $\vec{V} \sim V_m\hat{r}$, and $\vec{\nabla}\cdot\vec{V} \sim 2V_m/r$, both *GDMF* terms, $\vec{\nabla}[\rho(\vec{\nabla}\cdot\vec{V})]$ and $\vec{\nabla}[\vec{V}\cdot(\vec{\nabla}\rho)]$, consistently point to the source, in the $-\hat{r}$ direction.

Finally, consider the second velocity profile scenario (Section 9.2), using an accelerated radial velocity profile in combination with a Gaussian radial density profile. Since the "small r" approximate scalings are $\rho \sim \rho_o[1 - (r/a)^2]$,

$\partial\rho/\partial r \sim -2\rho_o r/a^2$, $\vec{V} \sim V_m(r/b)\hat{r}$, and $\vec{\nabla}\cdot\vec{V} \sim (3V_m/b)(1-2r/3b)$, both $\overrightarrow{GDMF}$ terms once again consistently point to the source, in the $-\hat{r}$ direction.

10. A theoretical analysis of fluxotaxis and chemotaxis in 3D

The previous section focused primarily on the merits of fluxotaxis in the context of sources and sinks. In this section, we address the question of how well fluxotaxis performs in a 3D environment in the absence of sinks and sources other than the one source – the emitter. We assume here that diffusion dominates and advection plays a peripheral and subsidiary role. Since chemotaxis is the most popular CPT method designed to address diffusion, we analyze both fluxotaxis and chemotaxis with a focus on chemical diffusion.

To emphasize diffusion, here and in all of the appendices the flow speed is assumed to be temporally invariant, where the velocity field has a constant radial component. The derivation in this section decomposes the density ρ into its 3D components, and delves deeper into the analysis that compares fluxotaxis with chemotaxis assuming two specific, but common types of emitters. Since we focus more on the chemical and less on the advective flow, we further refine our notion of the chemical density, ρ. In particular, it is useful to consider a particle number density formulation, $n(\vec{r},t)$, sometimes written as simply n, at a 3D spatial coordinate $\vec{r} = x\hat{i} + y\hat{j} + z\hat{k}$, and time t. The particle number density n is the chemical particle count per unit volume. Then, the mass density can be split up into its components, i.e. $\rho = mn$, which is the mass m of each chemical particle times the number density n. Therefore, our prior expression for the mass flux, $\rho\vec{V}$, can be rewritten as $mn\vec{V}$. Likewise, the 3D density diffusion equation can be written in terms of n as:

$$\frac{\partial n(\vec{r},t)}{\partial t} = D\nabla^2 n(\vec{r},t). \tag{8}$$

This equation describes the evolution of the density of particles, $n(\vec{r},t)$, as a function of the spatial coordinate, $\vec{r}$, and time, t. It is particularly useful to consider a temporal initial condition ($t = 0$) where all the N particles are contained at the origin ($\vec{r} = \vec{0}$), represented as a Dirac delta function particle source, where $n(\vec{r},0) = N\delta^3(\vec{r}) = N\delta(x)\delta(y)\delta(z)$. In Section 6.2, we showed that the diffusion coefficient D is proportional to the square of the mean free path, λ_{mfp}, divided by the collision time, τ_{coll}, or $D \propto \lambda_{\text{mfp}}^2/\tau_{\text{coll}}$.

The conserved total number of particles (by the conservation of mass principle) is:

$$N = \int_{-\infty}^{\infty} n(\vec{r},t)\mathrm{d}^3\vec{r} \quad = \int_{-\infty}^{\infty} n(\vec{r},0)\mathrm{d}^3\vec{r} \quad = \int_{-\infty}^{\infty} N\delta^3(\vec{r})\mathrm{d}^3\vec{r}.$$

Using a Green's function technique, we can solve the diffusion equation (8) for any type of emitter (Appendix 1.1). A Green's function is a Delta function source solution of inhomogeneous partial differential equations with boundary conditions. Ultimately, the Green's function for the diffusion equation can be used by integrating over the emitter source of particles to achieve the density solution for any localized particle emitter source. Appendix 1.2 provides a detailed solution for a single-puff emitter, and 1.3 addresses the case for a continuously emitting constant-speed emitter. The following two subsections address each of these cases, in turn.

10.1 The single-puff emitter case

Appendix 1.2 shows that the solution to the 3D diffusion equation (8), with a delta function initial condition and a single-puff emitter, is the Gaussian evolution function:

$$n(\vec{r},t)=\frac{N}{(4\pi Dt)^{3/2}}\,e^{-r^2/4Dt}. \tag{9}$$

The interpretation of such diffusion of chemical concentration, where the density is initially peaked at the origin, $\vec{r}=\vec{0}$, is that after a time t, the density has spread over a radial distance $r\approx\sqrt{4Dt}$ from the origin, where the amplitude is down by an e-folding level $1/e$, and the width of the density distribution increases as the square root of time, or $\Delta r\propto\sqrt{t}$, as shown in Figure 9.

Given that we now have the solution to the diffusion equation (8) for the case of a single-puff emitter and Gaussian density profile, we next consider the question of whether the $\overrightarrow{GDMF}$ of fluxotaxis and the gradient of chemotaxis are good guides (predictors of the direction of the source) for CPT. The pertinent question is, "Do either or both of these formulas/algorithms point in the correct direction toward the source emitter?" This question was answered earlier in Section 9 for fluxotaxis in 1D. Here, we extend the analysis to 3D, refine it with the particle density number $n(\vec{r},t)$, and compare it to chemotaxis. The hope is that the $\overrightarrow{GDMF}$ of fluxotaxis will be as effective as the density gradient ($\vec{DG}$) of chemotaxis, where chemotaxis is considered to be the "golden standard" for performance in domains that are dominated by diffusion, which is what we consider in this section.

$\overrightarrow{GDMF}$ analysis applied to the solution equation (9) of the 3D diffusion equation (8) is achieved by considering $n(\vec{r},t)=Ae^{-Br^2}=n(\vec{r})$, where we dropped the time notation by letting $A=N/(4\pi Dt)^{3/2}$ and $B=1/(4Dt)$, which in turn allows us to write $\partial n/\partial r$ as $n'=-2ABre^{-Br^2}$. To simplify the derivation further, note that for this case we are interested in a stationary outward flow velocity profile, $\vec{V}=V_o\hat{r}$, where V_o is the magnitude of the velocity at $t=0$, so that $\vec{\nabla}\cdot\vec{V}=2V_o/r$. Keeping in mind that $\rho=mn$, the two components of the $\overrightarrow{GDMF}$ vector can be expressed as:

$$1.\ \vec{\nabla}[\rho(\vec{\nabla}\cdot\vec{V})]=\frac{\partial}{\partial r}\left(\frac{2V_oAm}{r}e^{-Br^2}\right)\hat{r}\quad=2V_oAm\left(\frac{-1}{r^2}-2B\right)e^{-Br^2}\hat{r}$$
$$=-2V_oAm\left(\frac{1}{r^2}+2B\right)e^{-Br^2}\hat{r}.$$

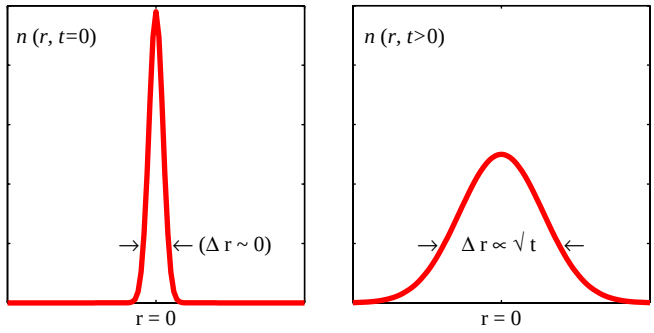

Figure 9.
Temporal growth of a chemical density profile generated by a single-puff source

$$2.\ \vec{\nabla}[\vec{V}\cdot(\vec{\nabla}\rho)] = -2ABm\frac{\partial}{\partial r}\left(V_o r e^{-Br^2}\right)\hat{r} = -2V_o ABm\left(1-2Br^2\right)e^{-Br^2}\hat{r}.$$

Substituting these values into the $\overrightarrow{GDMF}$ (equation (4)) yields:

$$\overrightarrow{GDMF} = m\vec{\nabla}\left[\vec{\nabla}\cdot(n\vec{V})\right] \quad = -2mV_o AB\left(3+\frac{1}{Br^2}-2Br^2\right)e^{-Br^2}\hat{r}.$$

For plume-tracing robots inside the inflection point $r_o = 1/\sqrt{2B}$ (i.e. the region where CPT is predictive), where:

$$n''(r_o) = -2AB\left(1-2Br_o^2\right)e^{-Br_o^2} = 0,$$

so that $n'' < 0$ for $0 \le r < r_o$, the $\overrightarrow{GDMF}$ vector correctly predicts the source direction to be:

$$\hat{s}\mathrm{GDMF} = \frac{\overrightarrow{GDMF}}{|\overrightarrow{GDMF}|} = -\hat{r}.$$

Further assurance is provided by the fact that in the limit:

$$\lim_{r\to 0}|\overrightarrow{GDMF}| \to \infty, \tag{10}$$

$$\lim_{r\to r_o}|\overrightarrow{GDMF}| \to 8mV_o AB > 0, \tag{11}$$

indicating that the $\overrightarrow{GDMF}$ vector is an effective and accurate source direction predictor for all of the relevant radial domains.

Next, consider a chemotaxis-driven robot, which tracks the density gradient $\vec{DG} = \vec{\nabla}n$. For the same Gaussian density diffusion function:

$$\vec{DG} = \vec{\nabla}\left[Ae^{-Br^2}\right] \quad = \frac{\partial}{\partial r}\left[Ae^{-Br^2}\right]\hat{r} \quad = -2ABre^{-Br^2}\hat{r}.$$

While this simpler scheme also predicts the correct source direction:

$$\hat{s}DG = \frac{\vec{DG}}{|\vec{DG}|} = -\hat{r},$$

the problem with chemotaxis-guided search is that for the radial flow configuration, where the chemical density profile contains a maximum point at $\vec{r} = \vec{0}$ (Figure 9), the magnitude DG of DG goes to zero when $\vec{r} \to \vec{0}$, i.e.:

$$\lim_{\vec{r}\to\vec{0}}|\vec{DG}| = \lim_{r\to 0} 2ABre^{-Br^2} \quad = 0.$$

This means that as chemotaxis-operated vehicles approach the source, their navigation vector $\vec{DG}$ disappears. In contrast to this, the quality of the $\overrightarrow{GDMF}$ prediction is maximized at the source location, as shown in (10), so that the magnitude of $\overrightarrow{GDMF}$ is highest near the origin of the plume, since $|\overrightarrow{GDMF}| \to MAX$ as $\vec{r} \to \vec{0}$, and the predictive capability of $\overrightarrow{GDMF}$ increases as fluxotaxis-controlled vehicles approach the

emitter. In practice, an important additional benefit of the fluxotaxis approach is that the plume-tracing region where $\overrightarrow{GDMF}$ correctly finds the emitter grows in time as $\Delta r \propto \sqrt{t}$.

The conclusion from this analysis is that there is a marked contrast between fluxotaxis and the more naïve density gradient-following chemotaxis scheme in the case of a single-puff emitter. Although both algorithms point in the correct source direction, as the robots get closer to the source emitter, chemotaxis loses its predictive ability, but the predictive ability of fluxotaxis improves. This is because chemotaxis uses only the first derivatives of $n(\vec{r}, t)$ as guidance, whereas fluxotaxis also uses the second derivative. Therefore, in the case of a single-puff emitter, fluxotaxis is actually preferable to chemotaxis – despite the fact that the two are being compared under diffusion conditions that are most favorable for chemotaxis!

10.2 The continuous emitter case

So far, this paper has focused on Gaussian density profiles. To demonstrate the generality of our approach, this section extends our analysis to handle the case of a continuous emitter, where the chemical density profile is an inverse radius decay function.

From Appendix 1.3, which solves the diffusion equation (8) for the continuous emitter case, we see that in the limiting case as $t \rightarrow \infty$, the particle density solution to the 3D diffusion equation is:

$$n(\vec{r}, \infty) = \frac{S_o}{4\pi D r}, \tag{12}$$

which has an inverse radius decay (i.e. $n \sim 1/r$), so that the number of particles drops off as the inverse of the radius from the source emitter. This limiting case solution for $t \rightarrow \infty$ is called the steady state solution. The steady state solution is used here for simplicity of exposition, though the conclusions presented in this section hold prior to the steady state as well. Figure 10 shows a graph of this inverse radius decay function. Here, we are assuming that S_o is the initial particle source rate (number of particles per unit time) from the source emitter, and $S(t) = S_o$ for all time, t, at the origin $\vec{r} = \vec{0}$, i.e. a constant rate emitter. Also, the total number of particles N increases with time as $N = tS_o$. Furthermore, this solution assumes a constant outward velocity profile, i.e. $\vec{V} = V_o\hat{r}$, where $\hat{r}$ is a unit vector pointing away from the emitter, and V_o is the initial magnitude of the chemical velocity at $t = 0$.

Given that we now have the steady state solution (12) to the diffusion equation (8) for the case of a continuous emitter, we next consider the question of whether the $\overrightarrow{GDMF}$ of fluxotaxis and the gradient of chemotaxis are good guides (predictors of the direction of the source) for CPT, i.e. "Do either or both of these formulas/algorithms point in the correct direction toward the source emitter?" in a continuous source situation dominated by diffusion.

To calculate the $\overrightarrow{GDMF}$ used by fluxotaxis, and the $\vec{DG}$ used by chemotaxis, we first need to find the spatial derivatives of $n(\vec{r}, t)$. After some algebraic manipulation we find that:

$$n' = \frac{-S_o r^{-2}}{4\pi D}$$

and:

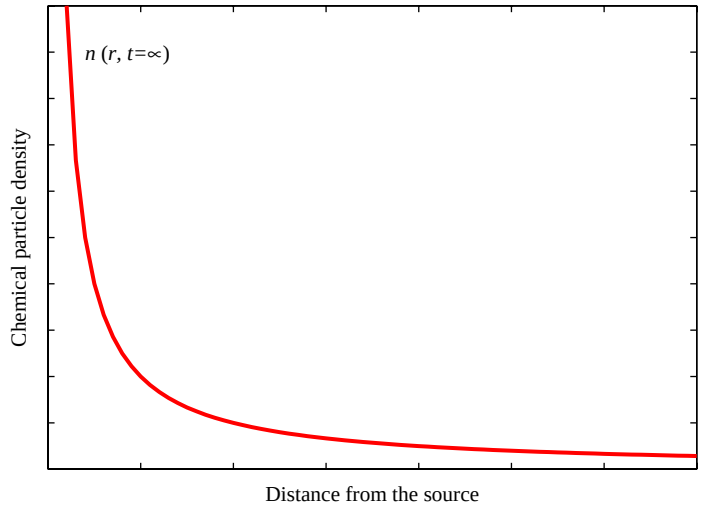

Figure 10.
Steady-state chemical concentration profile as a function of distance from the center of a continuous chemical emitter, expressed in terms of the particle density function $n(\vec{r}, t)$

$$n'' = \frac{S_o r^{-3}}{2\pi D}.$$

We also need further calculations to find $\overrightarrow{GDMF}$. Recall from equation (4) that the $\overrightarrow{GDMF}$ is composed of two terms:

(1) $\vec{\nabla}\left[\rho(\vec{\nabla}\cdot\vec{V})\right]$, which is the gradient of the product of density and the divergence of the velocity field; and

(2) $\vec{\nabla}\left[\vec{V}\cdot(\vec{\nabla}\rho)\right]$, which is the gradient of the flow velocity field in the direction of the chemical gradient.

In spherical coordinates, where $n = n(\vec{r}, t)$ and $\vec{V} = V_o\hat{r}$, the needed gradient and divergence results are: $\vec{\nabla}n = (\partial n/\partial r)\hat{r}$ and $\vec{\nabla}\cdot\vec{V} = r^{-2}(\partial/\partial r)(r^2 V_o)$. Note that $\vec{\nabla}\cdot\vec{V} = 2V_o/r$. Expanding, we get:

$$1.\ \vec{\nabla}\left[\rho(\vec{\nabla}\cdot\vec{V})\right] = \frac{2mV_oS_o}{4\pi D}\ \frac{\partial}{\partial r}(r^{-2})\hat{r} \quad = -\frac{mV_oS_o}{\pi D}\ \frac{1}{r^3}\hat{r}.$$

$$2.\ \vec{\nabla}\left[\vec{V}\cdot(\vec{\nabla}\rho)\right] = -\frac{mV_oS_o}{4\pi D}\ \frac{\partial}{\partial r}(r^{-2})\hat{r} \quad = \frac{mV_oS_o}{2\pi D}\ \frac{1}{r^3}\hat{r}.$$

Summing these two terms, we find that the $\overrightarrow{GDMF}$ is equal to:

$$\overrightarrow{GDMF} = -\frac{mV_oS_o}{2\pi Dr^3}\hat{r}. \tag{13}$$

This $\overrightarrow{GDMF}$ vector points in the direction:

$$\hat{s}\mathrm{GDMF} = \frac{\overrightarrow{GDMF}}{|\overrightarrow{GDMF}|} = -\hat{r}.$$

In other words, the $\overrightarrow{GDMF}$ points toward the origin in a radially symmetric stationary flow. This is exactly as desired.

Now, consider the density gradient, $\vec{DG}$, approach used by chemotaxis:

$$\vec{DG} = n'\hat{r} \quad = -\frac{S_o}{4\pi Dr^2}\hat{r}. \tag{14}$$

We can see that chemotaxis also predicts the correct source direction since:

$$\hat{s}\mathrm{DG} = \frac{\vec{DG}}{|\vec{DG}|} = -\hat{r}.$$

The conclusion is that both the $\overrightarrow{GDMF}$ and the $\vec{DG}$ vectors point in the correct source direction, and therefore fluxotaxis is as effective as chemotaxis for CPT domains where diffusion dominates the spread of the chemical away from a continuous emitter. This is a very reassuring result.

In addition to pointing toward the source emitter, both $\overrightarrow{GDMF}$ and $\vec{DG}$ have another advantage in the case of a continuous chemical source: they maximize their values at the emitter[6], $\vec{r} = \vec{0}$, so that $|\overrightarrow{GDMF}| \rightarrow MAX$. Thus, the predictive capabilities of $\overrightarrow{GDMF}$ and $\vec{DG}$ are both maximized close to the source emitter. However, it should be noted from equations (13) and (14) that as $\vec{r} \rightarrow \vec{0}$, the magnitude of the fluxotaxis predictor, $|\overrightarrow{GDMF}| \propto r^{-3}$, becomes larger than that of the chemotaxis predictor, $|\vec{DG}| \propto r^{-2}$.

11. Experimental confirmation of the theory

Our theory predicts that fluxotaxis should be an exceptionally adept CPT strategy, because it is based on fluid dynamics principles of source emitters. Now that we have fully developed our theory, let us consider how these results transfer to realistic 2D simulated chemical plume scenarios. In particular, how do the three CPT strategies (chemotaxis, anemotaxis and fluxotaxis) perform on plumes within varying flow regimes, such as laminar and turbulent? How do they perform under realistic conditions such as obstacles of varying number and size, anisotropic emitters, and unpredictable local variations in chemical density?

To address these questions, we have explored the performance of the three CPT strategies in the context of a "standard" CPT simulator that allows methodical perturbations of the emitter, flow, and simulated robot swarm characteristics. This "standard" CPT simulator was developed by Farrell *et al.* (2002). It is specifically designed for running large numbers of computationally efficient CPT experiments. The instantaneous and time-averaged results, as well as the multi-scale properties, of this simulator match those of real-world chemical plumes. Essentially, chemical puffs are emitted from a stationary chemical source emitter, and are carried advectively by the ambient wind flow. The regularity and frequency of puff emissions imply that this simulation more accurately models how a continuous, rather than a single-puff, emitter would act in the real world. Both advection and diffusion are modeled. The airborne emission is typically anisotropic, and it "meanders" dynamically as the wind currents

change over time. We have extended Farrell *et al.*'s original simulator to model obstacles, using fluid boundary conditions.

Many experiments have been run and there is only room here for a brief qualitative summary of them. For all of the quantitative and other specifics (Zarzhitsky *et al.*, 2004a, 2005; Zarzhitsky, 2008). We have varied the fluid flow to range from laminar to turbulent, varied the number and sizes of obstacles, the number and initial locations (e.g. initial distance from the emitter) of the simulated robots, and the inter-robot distances, and we have defined a variety of practical performance metrics. The latter include the frequency of emitter enclosure and the time of first detection of the emitter. Realistic robot constraints are modeled, such as realistic dimensions, maximum acceleration, sensor recovery time, and sensor and communication ranges. Also, both synchronous and asynchronous (multi-threaded) versions of the simulation have been implemented and tested, with no noticeable differences in the results.

Our simulation results demonstrate that the performance of fluxotaxis is outstanding, and it consistently outperforms chemotaxis and anemotaxis, based on all of our performance metrics (Zarzhitsky *et al.*, 2004a, 2005; Zarzhitsky, 2008), regardless of the perturbations in the flow regime, number and sizes of obstacles, and all of the other independent parameter variations. Chemotaxis is frequently fooled by local density maxima, especially when the environment has obstacles, thereby giving it poor "first arrival time" results. Anemotaxis often overshoots the emitter, thereby giving it poor "detection frequency" results. Fluxotaxis outperforms both chemotaxis and anemotaxis on both of these metrics[7].

In some of the experiments, the swarm size is increased methodically. The results demonstrate that CPT is inherently a swarm application. The bigger the swarm, the more effective and efficient the robots are at CPT.

Many of these results have been further confirmed in two other simulations, namely, a simulation of an extremely challenging urban environment containing large closely-spaced buildings of different sizes and shapes that substantially occlude the chemical flow (Agassounon *et al.*, 2009), and in a simulation of a deep ocean environment. For the latter, the CPT-driven robots are required to locate a hydrothermal vent, and there are sinks in the environment to fool the CPT on its way toward the source vent (Spears, 2009). Chemotaxis (which tracks heat, salinity or turbidity gradients) consistently gets trapped by the sinks; anemotaxis does not get trapped but it overshoots the hydrothermal vent. On the other hand, fluxotaxis flees the sinks and homes right in on the vent source, quickly, accurately, and consistently. This is excellent confirmation of our theoretical results in this paper.

Finally, recall that our theory in the previous Section 10.2 predicts that as the robots get closer to a continuous emitter, the magnitude of the fluxotaxis predictor will become larger than that of the chemotaxis predictor. When the speed of the robots is directly proportional to the magnitude of the virtual attractive force on them (toward the emitter), confirmation of this theoretical prediction should come in the form of faster convergence of the robot swarm upon a nearby emitter when using fluxotaxis than when using chemotaxis. We have in fact seen experimental confirmation. In particular, when watching the simulation it can be seen that when a $\overrightarrow{GDMF}$-driven fluxotaxis group of robots is near a source, it very rapidly "snaps" right onto the emitter location, whereas a chemotaxis-driven group of robots slowly moves toward the nearby emitter.

12. Summary and future work

In this paper, we presented an analysis of the CPT problem in the context of the principles of fluid dynamics. Our fluid-based physicomimetic approach is posed as a viable and important alternative to the biomimetic approaches outlined in Cowen (2002). The main focus of our research is a physics-based chemical source localization algorithm that can be implemented on a swarm of autonomous cooperating mobile robots. Our interest in the physicomimetic engineering of intelligent multi-robot systems stems from the relative ease with which such designs and algorithms can be analyzed and perfected. This approach allows us to apply similar mathematical techniques for both the design and analysis of our robot systems. In particular, the field of fluid dynamics provides numerous theoretical and experimental techniques that yield a great deal of insight into the behavior of chemical flows. It is therefore intuitive to apply a physicomimetic solution to a problem that has such a strong physical foundation.

The most immediate and practical contribution of this work is the development of a robust, theory-based chemical source identification procedure. We have shown how the divergence of chemical mass flux can be used to identify any chemical emitter. Our source declaration method can be readily combined with any of the popular CPT algorithms, such as chemotaxis or anemotaxis, to significantly improve the emitter localization accuracy with minimal changes to the existing design. The precision of our source identification method increases with the number of simultaneous observations of the chemical plume, which implies that the CPT performance can be improved substantially by increasing the number of plume-tracing robots in the swarm, thereby realizing the benefit of cooperation between multiple vehicles. The decentralized nature of our solution translates into unrestricted scalability and excellent resiliency of the physics-based algorithm.

In addition to the development of the source identification procedure, we constructed a complete plume-tracing algorithm called fluxotaxis to take advantage of the information contained within the chemical mass flux. We showed that fluxotaxis, having been developed from the fundamental principles of fluid dynamics, captures the salient features of both chemotaxis and anemotaxis – the two CPT algorithms that have been derived via emulation of biological systems. By relying on standard theoretical fluid-dynamics analysis techniques, we show in this paper that a CPT system that follows the gradient of the DMF performs at least as well as the chemotaxis technique in the presence of a continuous chemical source, and it actually outperforms chemotaxis on plumes produced by single-puff emitters. Complementary experimental results showing the superior performance of fluxotaxis over chemotaxis and anemotaxis on a wide variety of realistic laminar and turbulent chemical flow regimes, with obstacles present, are summarized here.

The experimental confirmation of our theoretical results (Section 11) demonstrates that the application of the physicomimetic approach for swarm design facilitates the development of CPT algorithms whose behavior is predictable. The ability to set bounds on the anticipated task performance in a theory-guided manner opens the door to the possibility of formalizing physicomimetic systems, thus widening the range of potential applications to include mission-critical tasks.

In conclusion, the theoretical framework we have established in this paper has been instrumental in allowing us to develop our algorithms further into a functioning CPT system, which we have implemented and tested under a variety of real-world

conditions. In the future, we intend to test and compare CPT algorithms in a variety of sophisticated simulators (Iacono and Reynolds, 2008), including 3D simulators. We also plan to compare the algorithms on observed real chemical plume data, such as the data that we have collected from our University of Wyoming Distributed Robotics Laboratory CPT chemical/wind tunnel, and the plume data collected by others (Crimaldi and Koseff, 2006).

Furthermore, we plan to compare CPT algorithms on our team of laboratory robots. Currently, the robots have chemical sensors, and the design and integration of miniature anemometer sensors are in progress. Once these sensors have been integrated, the next step will be to test and compare all of the CPT algorithms, including fluxotaxis, on the robots under varying indoor and outdoor conditions. Our ultimate objective is to understand the pros and cons of each of these strategies, and to clearly identify their regions of expertise in the space of possible environments.

Notes

1. Multiple emitters induce a separation of the swarm into smaller sub-swarms (Zarzhitsky and Spears, 2005).
2. For an excellent and comprehensive survey of CPT approaches that is much more general than the one herein, and which includes a categorization of fluxotaxis as a swarm-based strategy (Kowadlo and Russell, 2008).
3. The 1D results in this section originally appeared in Zarzhitsky *et al.* (2004c).
4. Without an explicit global reference frame, robots can share their local coordinate systems and establish a common frame of reference using standard transformations (Craig, 1989).
5. We use the expansion "small r" to mean that the lowest (nonzero) order Taylor series expansions have been performed on all density and velocity terms.
6. If the velocity is discontinuous at $\vec{r} = \vec{0}$, we can alternatively get the same result with $\lim_{\vec{r} \to \vec{0}}$.
7. Note that in general, any of the three CPT algorithms can get stuck at a local maximum that has features of a fluid "source" but is not the true overall global source emitter. Without a global perspective of the fluid field, or an infinite-sized swarm, this is a potential problem with any CPT strategy. A proper setting of the time period τ can help fluxotaxis to overcome this problem (Section 5).

References

Agassounon, W., Spears, W., Welsh, R., Zarzhitsky, D. and Spears, D. (2009), "Toxic plume source localization in urban environments using collaborating robots", paper presented at IEEE Conference on Technologies for Homeland Security.

Anderson, J.D. (1995), *Computational Fluid Dynamics*, McGraw-Hill, New York, NY.

Balkovsky, E. and Shraiman, B. (2002), "Olfactory search at high Reynolds number", *PNAS*, Vol. 99 No. 20, pp. 12589-93.

Board on Atmospheric Sciences and Climate (2003), *Tracking and Predicting the Atmospheric Dispersion of Hazardous Material Releases: Implications for Homeland Security*, National Academies Press, Washington, DC.

Buscemi, L., Prati, M. and Sandini, G. (1994), "Cellular robotics: behaviour in polluted environments", paper presented at 2nd International Symposium on Distributed Autonomous Robotic Systems, Springer-Verlag, New York, NY.

Caldwell, S.L., D'Agostino, D.M., McGeary, R.A., Purdy, H.L., Schwartz, M.J., Weeter, G.K. and Wyrsch, R.J. (1997), "Combating terrorism: federal agencies' efforts to implement national policy and strategy", Congressional Report GAO/NSIAD-97-254, Washington, DC, September 26.

Chemical Hazards and Poisons Division, HPA (2006), Chemical Hazards and Poisons Report, Chemical Hazards and Poisons Division, Health Protection Agency, London.

Coirier, W.J., Fricker, D.M., Furmanczyk, M. and Kim, S. (2005), "A computational fluid dynamics approach for urban area transport and dispersion modeling", *Env. Fluid Mech.*, Vol. 5 No. 5, pp. 443-79.

Cordesman, A.H. (2001), *Defending America: Asymmetric and Terrorist Attacks with Chemical Weapons*, Center for Strategic and International Studies (CSIS), Washington, DC.

Cowen, E. (2002), *Env. Fluid Mech.*, Vol. 2 Nos 1-2, special issue on Chemical Plume Tracing.

Craig, J.J. (1989), *Introduction to Robotics: Mechanics and Control*, Addison-Wesley, Reading, MA.

Crimaldi, J.P. and Koseff, J.R. (2006), "Structure of turbulent plumes from a momentumless source in a smooth bed", *Env. Fluid Mech.*, Vol. 6 No. 6, pp. 573-92.

Crimaldi, J.P., Koehl, M.A.R. and Koseff, J.R. (2002), "Effects of the resolution and kinematics of olfactory appendages on the interception of chemical signals in a turbulent odor plume", *Env. Fluid Mech.*, Vol. 2, pp. 35-63 (special issue on Chemical Plume Tracing).

Cui, X., Hardin, C.T., Ragade, R.K. and Elmaghraby, A.S. (2004), "A swarm approach for emission sources localization", *Proceedings of the 16th IEEE International Conference on Tools with Artificial Intelligence*, IEEE Press, Piscataway, NJ, pp. 424-30.

Decuyper, J. and Keymeulen, D. (1991), "A reactive robot navigation system based on a fluid dynamics metaphor", *Lecture Notes in Computer Science – Parallel Problem Solving from Nature (PPSN I)*, Vol. 496, Springer, Berlin, pp. 356-62.

Farrell, J.A., Murlis, J., Li, W. and Carde, R.T. (2002), "Filament-based atmospheric dispersion model to achieve short time-scale structure of odor plumes", *Env. Fluid Mech.*, Vol. 2 Nos 1/2, pp. 143-69 (special issue on Chemical Plume Tracing).

Feddema, J.T., Robinett, R.D. and Byrne, R.H. (2003), "An optimization approach to distributed controls of multiple robot vehicles", *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'03)*, IEEE Press, Piscataway, NJ.

Grasso, F. (2001), "Invertebrate-inspired sensory-motor systems and autonomous, olfactory-guided exploration", *Biol. Bull.*, Vol. 200, pp. 160-8.

Grasso, F. and Atema, J. (2002), "Integration of flow and chemical sensing for guidance of autonomous marine robots in turbulent flows", *Env. Fluid Mech.*, Vol. 2 Nos 1/2, pp. 95-114, special issue on Chemical Plume Tracing.

Hayes, A., Martinoli, A. and Goodman, R. (2001), "Swarm robotic odor localization", *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'01), Wailea*.

Hsu, S.S. (2003), "Sensors may track terror's fallout", *Washington Post*, June 2, p. A01.

Hughes-Hallett, D., Gleason, A.M. and McCallum, W. (1998), *Calculus, Single and Multivariable*, Wiley, New York, NY.

Iacono, G.L. and Reynolds, A.M. (2008), "Modelling of concentrations along a moving observer in an inhomogeneous plume. Biological application: model of odour-mediated insect flights", *Env. Fluid Mech.*, Vol. 8 No. 2, pp. 147-68.

Ishida, H., Tanaka, H., Taniguchi, H. and Moriizumi, T. (2006), "Mobile robot navigation using vision and olfaction to search for a gas/odor source", *Auton. Robots*, Vol. 20 No. 3, pp. 231-8 (special issue on Mobile Robot Olfaction).

Ishida, H., Nakamoto, T., Moriizumi, T., Kikas, T. and Janata, J. (2001), "Plume-tracking robots: a new application of chemical sensors", *Biol. Bull.*, Vol. 200, pp. 222-6.

Jeremić, A. and Nehorai, A. (1998), "Design of chemical sensor arrays for monitoring disposal sites on the ocean floor", *IEEE J. Ocean Eng.*, Vol. 23 No. 4, pp. 334-43.

Justus, K., Murlis, J., Jones, C. and Carde, R. (2002), "Measurement of odor-plume structure in a wind tunnel using a photoionization detector and a tracer gas", *Env. Fluid Mech.*, Vol. 21 Nos 1/2, pp. 115-42 (special issue on Chemical Plume Tracing).

Kazadi, S., Goodman, R., Tskikata, D., Green, D. and Lin, H. (2000), "An autonomous water vapor plume tracking robot using passive resistance polymer sensors", *Auton. Robots*, Vol. 9 No. 2, pp. 175-88.

Keymeulen, D. and Decuyper, J. (1994a), "The fluid dynamics applied to mobile robot motion: the stream field method", *Proceedings of the 1994 International Conference on Robotics and Automation (ICRA'94)*, Vol. 4, IEEE Press, Piscataway, NJ, pp. 378-85.

Keymeulen, D. and Decuyper, J. (1994b), "The stream field method applied to mobile robot navigation: a topological perspective", in Cohn, A.G. (Ed.), *Proceedings of the Eleventh European Conference on Artificial Intelligence (ECAI'94)*, Wiley, Chichester, pp. 699-703.

Kikas, T., Ishida, H., Webster, D. and Janata, J. (2001), "Chemical plume tracking. 1. Chemical information encoding", *Anal. Chemistry*, Vol. 73, pp. 3662-8.

Kowadlo, G. and Russell, R.A. (2008), "Robot odor localization: a taxonomy and survey", *Int. J. Robot. Res.*, Vol. 27 No. 8, pp. 869-94.

Krishnanand, K. and Ghose, D. (2006), "Glowworm swarm based optimization algorithm for multimodal functions with collective robotics applications", *Multiagent Grid Syst.*, Vol. 2 No. 3, pp. 209-22 (special issue on Recent Progress in Distributed Intelligence).

Liao, Q. and Cowen, E. (2002), "The information content of a scalar plume – a plume tracing perspective", *Env. Fluid Mech.*, Vol. 2 Nos 1/2, pp. 9-34 (special issue on Chemical Plume Tracing).

Lilienthal, A. and Duckett, T. (2003), "Approaches to gas source tracing and declaration by pure chemo-tropotaxis", *Proceedings of AMS, Karlsruhe, December 4-5*, pp. 161-71.

Lilienthal, A., Zell, A., Wandel, M. and Weimar, U. (2001), "Experiences using gas sensors on an autonomous mobile robot", *Proceedings of the 4th European Workshop on Advanced Mobile Robots (EUROBOT'01), Lund*, pp. 1-8.

Lilienthal, A., Ulmer, H., Frohlich, H., Stutzle, A., Werner, F. and Zell, A. (2004), "Gas source declaration with a mobile robot", *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'04), April 26-May 1, New Orleans, LA*, pp. 1430-5.

Lytridis, C., Kadar, E. and Virk, G. (2006), "A systematic approach to the problem of odour source localization", *Auton. Robots*, Vol. 20 No. 3, pp. 261-76, special issue on Mobile Robot Olfaction.

Marques, L., Nunes, U. and de Almeida, A. (2006), "Particle swarm-based olfactory guided search", *Auton. Robots*, Vol. 20 No. 3, pp. 277-87 (special issue on Mobile Robot Olfaction).

Parunak, H.V.D. and Brueckner, S. (2001), "Entropy and self-organization in multi-agent systems", *Proceedings of the International Conference on Autonomous Agents (AGENTS'01), Montreal*, pp. 124-30.

Russell, R.A., Kleeman, L. and Kennedy, S. (2000), "Using volatile chemicals to help locate targets in complex environments", *Proceedings of the Australian Conference on Robotics and Automation (ACRA'00), August 30-September 1, Melbourne*.

Sandini, G., Lucarini, G. and Varoli, M. (1993), "Gradient driven self-organizing systems", *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'93), July 26-30, Yokohama*.

Shinoda, T. (2003), "Koizumi's top-down leadership in the anti-terrorism legislation: the impact of political institutional changes", *SAIS Review*, Vol. 23 No. 1, pp. 19-34.

Song, Z. and Chen, Y. (2006), "Challenges and some results for the MAS-net project", *First Joint Emergency Preparation and Response Meeting*, pp. 258-64.

Spears, D. (2009), "Hydrothermal vent localization with a team of drones", University of Wyoming Technical Report.

Spears, D., Zarzhitsky, D. and Thayer, D. (2005), in Parker, L., Schultz, A. and Schneider, F. (Eds), *Multi-robot Systems: From Swarms to Intelligent Automata*, Vol. III, Springer-Verlag, New York, NY, pp. 211-22.

Spears, W. and Gordon, D. (1999), "Using artificial physics to control agents", *Proceedings of the IEEE Conference on Information, Intelligence, and Systems (ICIIS'99), Rockville, MD*, pp. 281-8.

Spears, W., Heil, R. and Zarzhitsky, D. (2005), "Artificial physics for mobile robot formations", *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (SMC'05)*, IEEE Press, Piscataway, NJ, pp. 2287-92.

Spears, W., Spears, D., Hamann, J. and Heil, R. (2004), "Distributed, physics-based control of swarms of vehicles", *Auton. Robots*, Vol. 17 Nos 2/3, pp. 137-62.

Spears, W., Hamann, J., Maxim, P., Kunkel, T., Zarzhitsky, D., Spears, D. and Karlsson, C. (2006), "Where are you?", *Proceedings of the SAB Swarm Robotics Workshop*, Springer-Verlag, New York, NY, pp. 129-43.

Thayer, D. (2008), "Theoretical foundations of chemical plume tracing with continuous and puff source emitters", University of Wyoming Technical Report.

Vergassola, M., Villermaux, E. and Shraiman, B. (2007), "Infotaxis as a strategy for searching without gradients", *Nature*, Vol. 445 No. 7126, pp. 406-9.

Weissburg, M., Dusenbery, D., Ishida, H., Janata, J., Keller, T., Roberts, P. and Webster, D. (2002), "A multidisciplinary study of spatial and temporal scales containing information in turbulent chemical plume tracking", *Env. Fluid Mech.*, Vol. 2 Nos 1/2, pp. 65-94, special issue on Chemical Plume Tracing.

Zarzhitsky, D. (2008), "A physics-based approach to chemical source localization using mobile robotic swarms", PhD thesis, University of Wyoming, Laramie, WY.

Zarzhitsky, D. and Spears, D. (2005), "Swarm approach to chemical source localization", *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (SMC'05)*, IEEE Press, Piscataway, NJ, pp. 1435-40.

Zarzhitsky, D., Spears, D. and Spears, W. (2004a), "Swarms for chemical plume tracing", *Proceedings of the IEEE Swarm Intelligence Symposium (SIS'04)*, IEEE Press, Piscataway, NJ, pp. 249-56.

Zarzhitsky, D., Spears, D. and Spears, W. (2005), "Distributed robots approach to chemical plume tracing", *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'05)*, IEEE Press, Piscataway, NJ, pp. 4034-9.

Zarzhitsky, D., Spears, D., Thayer, D. and Spears, W. (2004b), "A fluid dynamics approach to multi-robot chemical plume tracing", *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'04)*, IEEE Computer Society, Washington, DC, pp. 1476-7.

Zarzhitsky, D., Spears, D., Thayer, D. and Spears, W. (2004c), "Agent-based chemical plume tracing using fluid dynamics", *Lecture Notes in Computer Science*, Vol. 3328, Springer-Verlag, New York, NY, pp. 146-60.

Appendix 1. Supplementary derivations for the GDMF theory

1.1. General solution to the 3D diffusion equation

Here, we derive a general solution to the 3D diffusion equation for the density function, $n(\vec{r}, t)$:

$$\frac{\partial}{\partial t} n(\vec{r}, t) = D\nabla^2 n(\vec{r}, t),$$

where D is the diffusion coefficient that is defined in Section 6.2. We develop the needed Green's function as the density solution for a Dirac delta function initial condition ($t = 0$) of the chemical density at the origin $\vec{r} = \vec{0}$. The derivations assume time-independent velocity and a compressible fluid. $S(t)$ denotes the number of particles emitted per unit time from the source emitter which, for convenience, is assumed to be located at the origin of the coordinate system.

The density solution, $n(\vec{r}, t)$, of the 3D diffusion equation is obtained by integrating the product of the Green's function, $G(t, t')$, and the source $S(t')$, over time, t'. The Green's function is developed from the solution to the diffusion equation (assuming a delta function initial condition) using a time difference interval, $t - t'$. The Green's function is:

$$G(t, t') = \frac{1}{[4\pi D(t - t')]^{3/2}} e^{-(r^2/4D(t-t'))},$$

and the time-evolved solution to the 3D diffusion equation is:

$$n(\vec{r}, t) = \int_0^t G(t, t')S(t')\mathrm{d}t'. \tag{A1}$$

This Green's function propagates the source at every time t'. Here, it should be noted that the total number of particles emitted at the origin $\vec{r} = \vec{0}$, over a time t, is $N = \int_0^t S(t')\mathrm{d}t'$. This is because $\int_{-\infty}^{\infty} G(t, t')\mathrm{d}^3\vec{r} = 1$, or $\int_{-\infty}^{\infty}[n(\vec{r}, t)]\mathrm{d}^3\vec{r} = \int_0^t S(t')\mathrm{d}t'$, which is defined as N. Note that $\mathrm{d}^3\vec{r}$ is an abbreviation for $\mathrm{d}x\mathrm{d}y\mathrm{d}z$. Recall from calculus that $\int_{\infty}^{\infty}\mathrm{d}^3\vec{r} = \int_{\infty}^{\infty}\int_{\infty}^{\infty}\int_{\infty}^{\infty}\mathrm{d}x\mathrm{d}y\mathrm{d}z$.

The formula for $n(\vec{r}, t)$ in equation (A1) is the solution to the 3D density diffusion equation:

$$\frac{\partial}{\partial t} n(\vec{r}, t) = D\nabla^2 n(\vec{r}, t),$$

where $\nabla^2 n(\vec{r}, t) = r^{-2}\frac{\partial}{\partial r}(r^2\frac{\partial n}{\partial r})$, for any source, $S(t)$, at $\vec{r} = \vec{0}$. In addition, for a delta function source, $S(t) = N\delta(t)$, the usual density solution occurs, as in equation (9).

Note that $\lim_{t \to t'} G(t, t') = \delta^3(\vec{r})$ and $\frac{\partial}{\partial t} G(t, t') = D\nabla^2 G(t, t')$. Consequently, for a density source, $S(t)$, over the period of time $t' = 0$ to t, normalized to N, where $N = \int_0^t S(t')\mathrm{d}t'$, the 3D density result for all space, $\vec{r}$, is $n(\vec{r}, t) = \int_0^t G(t, t')S(t')\mathrm{d}t'$. The proof is as follows:

Proof:

$$\int_{-\infty}^{\infty} n(\vec{r}, t)\mathrm{d}^3\vec{r} = \int_0^t \left[\int_{-\infty}^{\infty} G(t, t')\mathrm{d}^3\vec{r}\right] S(t')\mathrm{d}t' \quad = \int_0^t S(t')\mathrm{d}t' \quad = N,$$

where $\int_{-\infty}^{\infty} G(t, t')\mathrm{d}^3\vec{r} = 1$. For $\vec{r} \neq \vec{0}$, we can rewrite:

$$\frac{\partial}{\partial t} n(\vec{r}, t) = \int_0^t \left[\frac{\partial}{\partial t} G(t, t')\right] S(t')\mathrm{d}t' D\nabla^2 n(\vec{r}, t) = \int_0^t [D\nabla^2 G(t, t')]S(t')\mathrm{d}t'.$$

Finally, with:

$$\frac{\partial}{\partial t} G(t,t') = D\nabla^2 G(t,t'),$$

we can conclude that:

$$\frac{\partial}{\partial t} n(\vec{r},t) = \int_0^t \left[D\nabla^2 G(t,t')\right] S(t')\mathrm{d}t' \quad = D\nabla^2 n(\vec{r},t). \quad \square$$

Thus, as previously stated, the density at all 3D spatial points $\vec{r}$ and time $t \geq 0$, driven by a source $S(t)$ at $\vec{r} = \vec{0}$, is correctly given by the integral:

$$n(\vec{r},t) = \int_0^t G(t,t')S(t')\mathrm{d}t'.$$

This general solution to the 3D diffusion equation is instantiated for the specific cases of single-puff and continuous emitters, in Appendices 1.2 and 1.3, respectively.

1.2. Solution to the 3D diffusion equation with a single-puff emitter

In this Appendix 1.2, we present the solution to the 3D diffusion equation that assumes a single-puff emitter. In other words, the emitter ejects a single chemical puff and then shuts down. We assume a delta function initial condition, where $S(t) = N\delta(t)$.

Thayer (2008) showed that the puff emitter solution to the 3D diffusion equation, with a delta function initial condition, is the Gaussian evolution function:

$$n(\vec{r},t) = \frac{N}{(4\pi Dt)^{3/2}} e^{-(r^2/4Dt)},$$

where $N = \int_{-\infty}^{\infty} n(\vec{r},t)\mathrm{d}^3\vec{r}$, and $n(\vec{r},0) = N\delta^3(\vec{r})$. We can verify that this is the value of N as follows[1]:

$$\begin{aligned}\int_{-\infty}^{\infty} n(\vec{r},t)\mathrm{d}^3\vec{r} &= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \frac{N}{(4\pi Dt)^{3/2}} e^{-(x^2/4Dt)} e^{-(y^2/4Dt)} e^{-(z^2/4Dt)} \mathrm{d}x\mathrm{d}y\mathrm{d}z \\ &= N\left[\frac{1}{\sqrt{4\pi Dt}}\int_{-\infty}^{\infty} e^{-(x^2/4Dt)}\mathrm{d}x\right]^3,\end{aligned}$$

due to spatial symmetry. Finally, substituting $x^2/4Dt = u^2$ so that $\mathrm{d}x = \sqrt{4Dt}\mathrm{d}u$, we obtain:

$$N\left[\frac{1}{\sqrt{4\pi Dt}}\int_{-\infty}^{\infty} e^{-(x^2/4Dt)}\mathrm{d}x\right]^3 = N\left[\frac{1}{\sqrt{\pi}}\int_{-\infty}^{\infty} e^{-u^2}\mathrm{d}u\right]^3 = N(1)^3 = N.$$

We can use this value of N to derive the initial condition solution $n(\vec{r},0)$. We now know that for all t, $N = \int_{-\infty}^{\infty} n(\vec{r},t)\mathrm{d}^3\vec{r}$. By taking the $t \to 0$ limit of the density function, $n(\vec{r},t)$, for the $|\vec{r}| = 0$ and $|\vec{r}| \neq 0$ cases, where:

$$\lim_{t\to 0} n(\vec{r},t) = \begin{cases} \infty, & |\vec{r}| = 0 \\ 0, & |\vec{r}| \neq 0, \end{cases}$$

and comparing this with the Dirac delta function definition, we can therefore conclude that $n(\vec{r},0) = N\delta^3(\vec{r})$. Note that this solution to the diffusion equation with a single-puff emitter is the usual Gaussian profile with spatial width proportional to $\sqrt{4Dt}$. See Figure 9 for a graph of this profile. For the verification that this is the solution to the diffusion equation in the case of a puff emitter, see the complete mathematical proof in Thayer (2008).

1.3. Solution of the 3D diffusion equation with a continuous emitter

For the case of a continuous chemical emitter, $S(t) = S_o$, so that a fixed number of particles per time step is being emitted at the $\vec{r} = \vec{0}$ source location. For this case, the solution to equation (8) is given by:

$$n(\vec{r}, t) = S_o \int_0^t G(t, t')\mathrm{d}t' \quad = S_o \int_0^t \frac{1}{[4\pi D(t - t')]^{3/2}} e^{-(r^2/4D(t-t'))}\mathrm{d}t'. \tag{A2}$$

Note that the total number of particles emitted from time 0 to time t is:

$$N = \int_{-\infty}^{\infty} n(\vec{r}, t)\mathrm{d}^3\vec{r} = S_o \int_0^t \mathrm{d}t' \int_{-\infty}^{\infty} G(t, t')\mathrm{d}^3\vec{r} = S_o \int_0^t (1)\mathrm{d}t' = tS_o.$$

In other words, $N = tS_o$.

To make the diffusion equation solution simpler, as needed in Section 10.2, we solve the integral in equation (A2). The solution to the Green's function integral is found as follows. Starting with a change of variables $(t - t') = \tau$ and $\mathrm{d}t' = -\mathrm{d}\tau$, we write:

$$n(\vec{r}, t) = S_o \int_0^t \frac{1}{(4\pi D\tau)^{3/2}} e^{-(r^2/4D\tau)}\mathrm{d}\tau.$$

Let $u = r/(\sqrt{4D\tau})$, $u^2 = r^2/(4D\tau)$, $\tau = r^2/(4Du^2)$, and $\mathrm{d}\tau = -(2r^2)/(4Du^3)\mathrm{d}u$, so that:

$$n(\vec{r}, t) = \frac{S_o}{\pi^{3/2}} \int_{r/\sqrt{4Dt}}^{\infty} \frac{r^2}{2Du^3} \frac{u^3}{r^3} e^{-u^2}\mathrm{d}u \quad = \frac{S_o}{4\pi Dr} \frac{2}{\sqrt{\pi}} \int_{r/\sqrt{4Dt}}^{\infty} e^{-u^2}\mathrm{d}u.$$

Note that this expression is related to the probability integral, or the error function *erf*, which is defined as $erf(\alpha) = 2/\sqrt{\pi}\int_0^{\alpha} e^{-u^2}\mathrm{d}u$. Here, $\alpha = r/\sqrt{4Dt}$, $erf(\infty) = 1$, and $erf(0) = 0$, while $(2/\sqrt{\pi})\int_0^{\alpha} e^{-u^2}\mathrm{d}u + (2/\sqrt{\pi})\int_{\alpha}^{\infty} e^{-u^2}\mathrm{d}u = 1$. Therefore, we have:

$$n(\vec{r}, t) = \frac{S_o}{4\pi Dr} \left\{1 - erf\left[\frac{r}{\sqrt{4Dt}}\right]\right\}.$$

Consider two limiting cases of this solution:

Case 1. For $t \to 0$, $erf(\infty) = 1$. This implies that $n(\vec{r}, 0) = 0$ for $|\vec{r}| > 0$.

Case 2. For $t \to \infty$, $erf(0) = 0$. This implies that $n(\vec{r}, \infty) = S_o/(4\pi Dr)$, which has an inverse radius decay, as noted in Section 10.2 and shown in Figure 10. In fact, it is this Case 2 solution that is addressed in the main body of the paper.

Finally, note that for small distances r and large times t, where $\alpha = r/\sqrt{4Dt} \ll 1$, the Taylor series expansion of the error function is useful:

$$erf(\alpha) = \frac{2}{\sqrt{\pi}} e^{-\alpha^2} \sum_{n=0}^{\infty} \frac{2^n}{1 \cdot 3 \cdots (2n + 1)} \alpha^{2n+1},$$

with:

$$n(\vec{r}, t) = \frac{S_o}{4\pi Dr} \{1 - erf(\alpha)\}.$$

In other words, we can solve for $n(\vec{r}, t)$ computationally, using this expansion, at any point in space and time – in order to calculate the GDMF. See Thayer (2008) for verification and a complete mathematical proof that this is indeed the solution to the diffusion equation in the case of a continuous emitter.

Note: 1. Recall from calculus that $\int_{\infty}^{\infty}\mathrm{d}^3\vec{r} = \int_{\infty}^{\infty}\int_{\infty}^{\infty}\int_{\infty}^{\infty}\mathrm{d}x\mathrm{d}y\mathrm{d}z$.

About the authors

Diana F. Spears received the PhD degree in Computer Science from the University of Maryland in 1990. She is currently an owner and senior research scientist at Swarmotics, LLC, and also an Adjunct Associate Professor in the Computer Science Department at the University of Wyoming (UW). Previously, she worked at NASA Goddard Space Flight Center, the National Institute of Standards and Technology, and as an Associate Professor of Computer Science at UW. She was an Associate Editor for the *Journal of Artificial Intelligence Research* and is currently on their Advisory Board. She is currently serving as an invited member on the AAAI Presidential Panel on the Future of Artificial Intelligence. She has over 80 publications, and is known internationally for her expertise on swarm robotics and multi-robot chemical plume tracing. Her other research interests include behaviorally assured adaptive and machine learning systems, and mathematical/graphical modeling and model reconstruction. Her personal web site is: www.cs.uwyo.edu/~dspears. Diana F. Spears is the corresponding author and can be contacted at: dspears@swarmotics.com

David R. Thayer received the PhD degree in Plasma Physics from MIT in 1983. He has been a faculty member of the Department of Physics and Astronomy at the UW since 2000. Previous positions have included: a senior scientist position at SAIC working directly under the CEO on projects including biotechnology, global change research, nanotechnology, submarine turbulent flows (DARPA funded), and fusion plasma transport (DOE funded), and a staff scientist position at Lawrence Berkeley Laboratory working on ICRF heating through turbulent edge plasma devices. Currently, in addition to CFD research with the Computer Science Department at UW, he works with the UW Mathematics and Electrical and Computer Engineering departments on non-thermal plasma reactor analysis for hydrogen production. He teaches many of the graduate and upper-division undergraduate physics courses at UW, including classical mechanics, plasma physics, mathematical and computational physics and quantum mechanics. He is a member of the American Physical Society.

Dimitri V. Zarzhitsky received the PhD degree in Computer Science from the UW in 2008. He is an active participant in several ongoing research efforts between government agencies and university laboratories, overseeing design, development, and deployment of unmanned systems for both civilian and military applications. Presently, he is working on improving effectiveness of cooperative sensing and control algorithms for distributed autonomous systems.

To purchase reprints of this article please e-mail: **reprints@emeraldinsight.com**
Or visit our web site for further details: **www.emeraldinsight.com/reprints**